UNIVERSITÀ VITA-SALUTE SAN RAFFAELE

CORSO DI DOTTORATO DI RICERCA
IN FILOSOFIA

CURRICULUM IN DISCIPLINE FILOSOFICHE

WHEN NEWS BREAKS THE NORMS OF JOURNALISM: AN INSTITUTIONAL APPROACH TO FAKE NEWS ABOUT CLIMATE CHANGE

Tutor: Prof.ssa Biancamaria Cepollaro



Tesi di DOTTORATO di RICERCA di Maurizio Mascitti

matr. 020165

Ciclo di dottorato XXXVII

SSD M-FIL/05 FILOSOFIA DEL LINGUAGGIO

Anno Accademico 2023/2024

CONSULTAZIONE TESI DI DOTTORATO DI RICERCA

Il/la sottoscritto/*I.* Maurizio Mascitti

Matricola / *registration number* 020165

nato a/ *born at* Ascoli Piceno

il/*on* 26/06/1997

autore della tesi di Dottorato di ricerca dal titolo / *author of the PhD Thesis titled*

When News Breaks the Norms of Journalism: An Institutional Approach to Fake News about Climate Change

☒ AUTORIZZA la Consultazione della tesi / AUTHORIZE *the public release of the thesis*

□ NON AUTORIZZA la Consultazione della tesi per sei mesi / DO NOT AUTHORIZE *the public release of the thesis for six months*

a partire dalla data di conseguimento del titolo e precisamente */ from the PhD thesis date, specifically*

Dal / *from* Al / *to*

Poiché /*because*:

□ l'intera ricerca o parti di essa sono potenzialmente soggette a brevettabilità/ *The whole project or part of it might be subject to patentability;*

□ ci sono parti di tesi che sono già state sottoposte a un editore o sono in attesa di pubblicazione/ *Parts of the thesis have been or are being submitted to a publisher or are in press;*

□ la tesi è finanziata da enti esterni che vantano dei diritti su di esse e sulla loro pubblicazione/ *the thesis project is financed by external bodies that have rights over it and on its publication.*

È fatto divieto di riprodurre, in tutto o in parte, quanto in essa contenuto / *Copyright the contents of the thesis in whole or in part is forbidden*

Data /Date 27/01/2025 Firma /Signature Maurizio Mascitti

DECLARATION

This thesis has been:

- composed by myself and has not been used in any previous application for a degree.
  Throughout the text I use both 'I' and 'We' interchangeably.
- has been written according to the editing guidelines approved by the University.

Permission to use images and other material covered by copyright has been sought and obtained. For the following image/s (specify), it was not possible to obtain permission and is/are therefore included in thesis under the "fair use" exception (Italian legislative Decree no. 68/2003).

All the results presented here were obtained by myself.

All sources of information are acknowledged by means of reference.

Acknowledgments

Over these three years of my PhD program, I have encountered many people along the way. With many of them, I had the pleasure of working together, while others enriched my experience in different ways. Each of them has held a unique significance for me.

First and foremost, I would like to thank Prof. Biancamaria Cepollaro, my supervisor. I am deeply grateful to her both for the immense help she provided throughout this journey and for supporting me during the most challenging moments of my PhD. If it hadn't been for her and the way she approached me, I might have given up the doctorate in the first year.

I would also like to thank Prof. Claudia Bianchi, both in her role as a temporary supervisor for my project and for believing in me and my work on multiple occasions during my time at UniSR. Lastly, I extend my gratitude to all the researchers and faculty members (past and present) at San Raffaele University from whom I received advice and support, either directly or indirectly, including, in particular, Francesca Boccuni, Laura Caponetto, Carlo Martini, and Raffaele Ariano.

But this milestone was also, if not primarily, made possible thanks to the presence of many colleagues who, over the years, became something more: friends. I would therefore like to thank them for making me feel less alone during the darkest moments and for enriching my PhD experience with stimulating conversations, laughter, and lightheartedness. I especially want to thank Luca Ausili, Alessandro Volpi, Eleonora Volta, Paolo Polizzi, Alessandro Volpe, Paolo Colizzi, Davide Battisti, Federico Bina, Francesca Cesarano, Giorgia Adorno, Giulio Pennacchioni, Simone Marsilio, Nicolò D'Agruma, Stefano Pinzan, Alessandro Anzà, Margherita Grassi, Martina Giovine and many others who, I hope, will forgive me for not including them in this endless list.

Then there are the professors and researchers I had the pleasure of meeting during my visiting period in Uppsala, to whom I owe a great deal. First and foremost, I would like to thank Prof. Jessica Pepp, who guided me throughout the six months I spent in the Department of Philosophy at Uppsala. I still fondly remember the long conversations in her office about my project, and I am profoundly indebted to her: she provided me with invaluable advice and insights for my research. Moreover, I must say that I have rarely encountered someone as kind and approachable as her.

In addition to her, I would also like to thank Delia Belleri for giving me the opportunity to present my research project, and Andreas Stokke for providing me with important perspectives through his critiques.

But in Sweden, I didn't just carry out research activities: I also met many people there who left a lasting impression on me. In particular, I would like to thank Jerry Faisal, my roommate, for making me feel at home by welcoming me into his circle of friends and introducing me to his wonderful companions (and for letting me enjoy so much delicious food!). I also want to thank Stefano Patriarca, with whom I formed a beautiful friendship that I hope will last forever.

Over these three years of my PhD, I have also received the support and companionship of many others, even from a distance. I would like to specifically thank Gabriele Giancola, Simone Conversano, Niccolò Prospero, Giosuè Di Girolamo, and Geremia Di Girolamo for the countless days we spent together in our 'office' on Discord. You have no idea how much your presence meant to me during the past year and a half, and I truly hope you will continue to be a part of my life journey for a long time to come.

Then there are my lifelong friends, who have been part of my journey even from afar, and my family, who never stopped supporting me throughout this rollercoaster of an adventure. A special thank you goes to my life partner, Michela, who, even during the many difficult moments when I faltered, gave me the love and support without which no undertaking would truly be possible – not even a PhD.

To Noemi,

who didn’t make it to see me reach this goal, but who would have so dearly wanted to.

Abstract

This dissertation redefines the philosophical concept of fake news to address an issue of conceptual necessity and to analyze cases of fake news related to climate change more effectively. These objectives are addressed in four chapters, with the first three focusing on redefining the concept and the final chapter applying it to climate change.

Chapter I explores the historical and definitional foundations of fake news, tracing its evolution and conceptual challenges. By examining conservative and privative accounts in epistemology, the chapter critiques their redundancy and highlights the necessity for a different approach.

Chapter II develops a novel understanding of news as an institutional entity, drawing on John Searle's theory of institutional facts. This perspective conceptualizes news as a status function stemming from journalistic practices, which confer specific rights and obligations on both producers and consumers. By framing news as an institutional construct, the chapter addresses the limitations of privative and conservative approaches, resolving ambiguities while emphasizing the normative expectations that distinguish news from other forms of communication.

Chapter III introduces a redefined notion of fake news as news that does not deserve its status, bifurcated into two complementary dimensions. The first (FN_A) pertains to violations of journalistic norms, such as failures in verification or transparency. The second (FN_B) reflects the failure of news to meet realistic consumer expectations informed by journalistic standards. By integrating these producer-and-consumer-focused perspectives, the chapter provides a comprehensive framework for understanding fake news as both a normative breach and a failure to fulfill public trust.

Chapter IV applies this redefinition to climate change, differentiating climate fake news from disinformation. It examines cases where climate news violates journalistic norms, such as deceptive headlines or articles exploiting audience biases. These cases reveal how fake news can reinforce pre-existing beliefs without directly disseminating falsehoods, employing a subtler strategy to influence perception. The chapter concludes by advocating for refined analytical tools to address the diverse notions of intentionality in disinformation phenomena, particularly those involving audience predispositions.

By anchoring fake news in journalistic norms and expectations, this thesis offers a framework for conceptual clarity and practical analysis, especially in contexts like climate change, where the term is frequently conflated with disinformation.

Abstract (italiano)

Questa tesi ridefinisce il concetto filosofico di fake news per affrontare un problema di necessità concettuale e analizzare in modo più efficace i casi di fake news legati al cambiamento climatico. Gli obiettivi vengono sviluppati in quattro capitoli: i primi tre sono dedicati alla ridefinizione del concetto, mentre l'ultimo lo applica al contesto del cambiamento climatico.

Il Capitolo I esamina le basi storiche e definitorie del concetto di fake news, ripercorrendo la sua evoluzione e le relative sfide concettuali. Attraverso l'analisi delle definizioni conservative e privative di fake news in epistemologia, il capitolo ne critica la ridondanza e sottolinea la necessità di un approccio diverso.

Il Capitolo II sviluppa una nuova concezione della notizia come entità istituzionale, ispirandosi alla teoria dei fatti istituzionali di John Searle. Da questa prospettiva, le notizie sono definite come una funzione di status che discende dalle pratiche giornalistiche, le quali attribuiscono specifici diritti e obblighi sia ai produttori che ai consumatori di notizie. Inquadrando le notizie come costrutti istituzionali, il capitolo affronta i limiti degli approcci conservativi e privativi, risolvendo alcune ambiguità e sottolineando le aspettative normative che distinguono le notizie da altre forme di comunicazione.

Il Capitolo III propone una nuova definizione di fake news come notizie che non meritano il loro status, articolata in due dimensioni complementari. La prima (FN_A) riguarda le violazioni delle norme giornalistiche, come la mancanza di verifica o trasparenza. La seconda (FN_B) si riferisce al mancato soddisfacimento delle aspettative realistiche dei consumatori, basate sugli standard giornalistici. Integrando queste prospettive, focalizzate rispettivamente sui produttori e sui consumatori, il capitolo offre un quadro completo per comprendere le fake news sia come violazione normativa sia come tradimento delle aspettative del pubblico.

Il Capitolo IV applica questa ridefinizione al cambiamento climatico, distinguendo le fake news climatiche dalla disinformazione. Analizza casi in cui le notizie sul clima violano le norme giornalistiche, ad esempio titoli ingannevoli o articoli che sfruttano i pregiudizi del pubblico. Questi esempi mostrano come le fake news possano rafforzare credenze preesistenti senza diffondere direttamente falsità, utilizzando strategie più

sottili per influenzare il pubblico. Il capitolo si conclude con un invito a sviluppare strumenti analitici più sofisticati per affrontare le diverse sfaccettature dell'intenzionalità nei fenomeni di disinformazione, in particolare quelli che coinvolgono le predisposizioni del pubblico.

Ancorando il concetto di fake news alle norme e alle aspettative giornalistiche, questa tesi intende offrire un quadro utile per una maggiore chiarezza concettuale e un'analisi pratica, soprattutto in contesti come il cambiamento climatico, in cui il termine "fake news" viene spesso confuso con la disinformazione.

Table of Contents

Introduction *1*

Chapter 1 ..... *7*

What is fake news: the state of the art *8*

1. Historical roots **8**

2. The definitional project and its challenges: Habgood-Coote's thesis **10**

2.1 Vs meaning instability 13

2.2 Vs conceptual necessity 14

2.3 Vs weaponization 15

2.4 Taking stock: how shall we think of fake news? 16

3. The current debate: conservative vs privative accounts **20**

4. The conservative view of fake news **22**

5. Doing without subjective conditions **28**

5.1 Deceiving the audience about content (A1) or source (A2) 29

5.2 Bullshitting the audience (B) 31

5.3 Disjunctive accounts (C) 34

5.4 The minimum audience condition 36

6. The problem of redundancy **38**

7. The privative view of fake news **41**

8. The imitation of genuine news: a critique **47**

8.1 What does empirical evidence say? 47

8.2 A forgery is (not) always a forgery 50

9. The problem of redundancy, again **55**

Conclusion **56**

Chapter 2 ..... *59*

The nature of news *60*

1. Searching for the culprit **60**

2. Prevailing positions on news **61**

3. News as reported information **63**

3.1 News is not just information 67

4. News as a status 72

5. What is an institution? 77

5.1 Collective intentionality 79

5.2 The assignment of function 80

5.3 Status functions 82

5.4 The general structure of institutional facts 83

5.5 Deontic powers 85

5.6 Summary 88

6. News as an institutional entity 88

6.1 Collective agreement and the value of news 96

6.2 On imitation of news appearance 98

Conclusion 104

Chapter 3 107

Fake news: an institutional account 108

1. The definition of fake news 108

2. Norms of journalism and realistic expectations 109

2.1 Typology of norms 113

2.2 Hierarchy of norms 121

3. The institutional account of fake news 126

3.1 The interpretation of “fake” in “fake news” 131

3.2 Properties of fake news 136

4. The pragmatics of fake news 143

4.1 Relevance Theory: staples 146

4.2 Clickbait headlines 155

4.3 Deceptive headlines 160

4.4 Targeted fake news 165

Conclusion 174

Chapter 4 177

Fake news about climate change: what’s new? 178

1. What is climate change? 178

2. Disinformation: the main enemy of climate science.....180

3. Is there a problem of climate fake news? 184

3.1 Climate fake news: what we know so far 186

4. Institutional climate fake news 192

4.1 Clickbait headlines on climate change 193

4.2 Deceptive headlines on climate change 197

4.3 Targeted climate fake news 200

Conclusion.....209

Concluding remarks 213

References 218

Introduction

"For the most part, we do not first see, and then define; we define first, and then see."

Walter Lippman, *Public Opinion*

Few challenges in human history rival the complexity and urgency of climate change. As one of the defining crises of our time, it is not merely an environmental issue but a societal, political, and moral one that demands collective action on an unprecedented scale. Climate change impacts every facet of life: ecosystems are destabilized, economies are strained, and vulnerable populations face existential threats. Despite the growing body of scientific evidence and the increasingly visible effects of global warming, the road to meaningful action remains riddled with obstacles.

Chief among these is the challenge of effectively communicating the gravity of the crisis to the public. A survey conducted in March 2022 by the Pew Research Center found that, overall, the majority of U.S. citizens described climate change as a major threat to the country's well-being. But this majority was fragile: just 54% of U.S. adults identified as seriously concerned, representing a slight decline since 2020. Furthermore, a stark political divide persists. Nearly eight-in-ten Democrats (78%) believed climate change poses a critical threat to the country's well-being; by contrast, only about one-in-four Republicans (23%) shared this view. Three years have passed since that survey, and now, considering the current political atmosphere in the United States amid Donald Trump's second term, the divide may have widened even further. In October 2024, the Pew Research Center found that an overwhelming majority of Democrats (91%) reported feeling sadness over climate news, whereas Republicans more often reacted with suspicion toward climate advocates, with 75% expressing such skepticism. The situation in the European Union, however, paints a more optimistic picture. The Special Eurobarometer No. 538, conducted between March and May 2023, revealed that approximately 77% of respondents across surveyed countries considered climate change a very serious problem, ranking its severity between 7 and 10 on a 10-point scale.

Over the past decade, scientists and commentators have frequently pointed to the role of disinformation in complicating the communication of the climate crisis. In the context of climate change, disinformation has been characterized by deliberate campaigns designed to mislead, often orchestrated by parties with vested interests. These campaigns include sensationalized reporting that obscures scientific nuance, deliberate dissemination of falsehoods, and viral content propagated on social media. The ramifications are profound: such campaigns distort public understanding, nurture skepticism about climate science, and erode trust in the scientific community. By undermining confidence, these false narratives hinder collective efforts to rally public support and implement the policies necessary to address the climate crisis. They create a fragmented information landscape in which verifiable scientific facts must compete with persuasive but misleading counter-narratives.

However, in recent years, traditional media outlets and even scientists themselves have expanded the ranks of enemies of climate science. Increasingly, the term “fake news” has entered the discussion as an additional, albeit related, challenge. While often used interchangeably with “disinformation”, the usage of “fake news” to describe climate misconceptions has gained momentum in public debate. Although not a neologism of the digital age, “fake news” has attracted considerable attention from the media and academia. Numerous studies have analyzed the sociology of fake news, its mimicry of legitimate journalism, and its primordial affinity with the digital environment. The phenomenon gained widespread attention during the 2016 U.S. presidential election, after which the term experienced a troubling semantic shift. Originally used to describe fabricated stories crafted to deceive, it increasingly became a partisan weapon, wielded primarily by pro-Trump supporters to discredit mainstream media outlets. In this capacity, “fake news” evolved into a pejorative term, disconnected from its original descriptive function.

Despite its semantic drift, the philosophical discourse on fake news – especially in epistemology – has sought to restore clarity and rigor to the concept. The overarching aim of philosophers in this domain has been to define “fake news” in a way that stabilizes its meaning and situates it within broader conceptual frameworks. A common strategy involves deconstructing the concept into more fundamental notions, such as truth and truthfulness. Many scholars emphasize two core elements of fake news: the

lack of truth, which may manifest as outright falsehoods or misleading implications, and the lack of truthfulness, reflecting either the producer's intent to deceive or their indifference to the truth. Others have underscored the parasitic nature of fake news, noting that it masquerades as legitimate journalism while subverting its norms.

A persistent issue in this literature is that prevailing definitions of fake news often fail to add significant explanatory value to established terms. For instance, defining fake news as false information deliberately spread to mislead closely mirrors the concept of disinformation. Similarly, when definitions omit the requirement of intent and describe fake news as false information spread without regard to truth, they align closely with the concept of misinformation. In both cases, the term risks redundancy, appearing as a conceptual duplicate that offers little unique insight. This redundancy problem undermines the theoretical necessity of the concept of fake news and fuels skepticism among philosophers who advocate for its abandonment. This is precisely the issue that motivates the sense and purpose of this dissertation.

The starting claim of this work is that we cannot fully understand the concept of fake news – or its role in shaping climate change discourse – without first examining the concept of news itself. News occupies a central role in contemporary society. It bridges individual awareness and collective reality, mediating how we interpret events, make decisions, and construct shared meaning. It sets the public agenda, influences policy debates, and provides the foundation for democratic engagement. Yet, in everyday discourse, news is often conflated with information, as if it were merely a conduit for transmitting facts from sender to receiver. This assumption is profoundly flawed, obscuring the complex processes through which news operates. To grasp the power and societal significance of news, it is essential to recognize a crucial distinction central to this work: information is an object – a semantic entity – while news is a status conferred upon information. This transformation occurs when information is produced according to the standards and norms of journalism, an institution that governs the selection, framing, and presentation of news. In this sense, news is not merely a vehicle for transmitting information; it is a socially constructed category endowed with authority and legitimacy through adherence to journalistic conventions. Any information emerging from a journalistic source acquires the status of news, irrespective of its content or context. These conventions – rooted in the ethics, practices, and traditions of

journalism – are what transform raw facts into narratives deemed newsworthy. The power of news lies in its ability to create legitimate expectations among consumers that the information they receive is accurate and significant. This status is ultimately conferred by public trust in the institution of journalism and its perceived legitimacy.

From these premises emerges the central proposal of this work: fake news is best understood as news that fails to deserve its status. At its core, fake news is not about the falsehood or inaccuracy of content. Rather, it represents a failure to uphold the institutional obligations of journalism – a breach of the professional commitments that underpin the practice. Fake news arises when information fails to meet the standards of journalism, either because it violates professional norms or falls short of the reasonable expectations of its audience. This dual perspective situates fake news as a phenomenon rooted in institutional trust and professional accountability, rather than merely the veracity of content.

The objective of this work is to develop a nuanced understanding of fake news by situating it within the institutional framework of journalism. By examining the interplay between journalistic norms, consumer expectations, and the status function of news, this dissertation sheds light on why fake news remains a valuable philosophical concept. It further distinguishes fake news from related phenomena such as misinformation and disinformation, arguing that fake news is not merely false news but a significant breakdown in journalistic practice. Importantly, this framework acknowledges that fake news can consist of factually accurate content that nevertheless fails to justify its status as news. By tying fake news to the institutional obligations of journalism, this approach avoids redundancy with other concepts and provides a clearer basis for understanding its unique characteristics.

A significant focus of this dissertation is also the pragmatics of fake news, emphasizing the role of norms and expectations in shaping journalistic communication. Unlike conversational exchanges, where mutual adjustment allows for clarification and feedback, journalism operates as a one-sided communicative act. Journalists must anticipate the needs of their audience and communicate with precision, clarity, and transparency to ensure comprehension. When these commitments are violated – intentionally or otherwise – the result is a breakdown in communication that undermines the quality of news. This breakdown is especially acute in the case of

climate fake news, where ideological divides amplify its impact on public discourse and hinder collective action. This brings us back to the specific context of fake news about climate change (or climate fake news).

The final portion of this dissertation applies the institutional account to climate fake news, demonstrating how this framework captures cases that elude traditional definitions rooted in truth and intentionality. Climate fake news often involves subtler violations of journalistic norms, such as sensationalist headlines, clickbait that prioritizes engagement over informativeness, or narratives designed to exploit ideological biases. These practices frequently target conservative or libertarian audiences predisposed to skepticism about anthropogenic climate change, leveraging their ideological commitments to magnify the impact of misleading narratives. By employing the institutional framework, this dissertation illuminates new dimensions of climate fake news, showcasing its value as a conceptual tool for understanding the interplay between journalism, trust, and news in the age of climate crisis.

Chapter 1

What is fake news: the state of the art

1. Historical roots

The term “fake news” has a long history behind it. Contrary to what most people believe, this expression did not intrude the public debate for the first time in 2016. At that time, many journalists and pundits were using “fake news” in the run-up to the US Presidential elections, when the existence of a campaign of information manipulation around elections led by Russian government was no longer just a sneaking suspicion (Mazzetti 2020). Later, the Collins Dictionary announced “fake news” as the Word of the Year for 2017, after its usage surged by 365% in the previous 12 months; thereof corroborating the novelty of the term, albeit involuntarily. According to the Collins Dictionary, fake news means «false, often sensational, information disseminated under the guise of news reporting». But there is plenty of evidence that such use of the term was already deployed before 2016, and sometimes much earlier.

Even though the phenomenon of fake news may have taken off at the same time that news began to circulate after Johannes Gutenberg invented the printing press in 1439 (Soll 2016), the term acquired general use in print since the end of the 19th century. For instance, a headline of *Cincinnati Commercial Tribune*, a major daily newspaper in Cincinnati (Ohio) now folded, reads “Secretary Brunnell Declares Fake News About His People is Being Telegraphed Over the Country”. The headline dates back to 7 June 1890. The term remained under the radar throughout the 20th century, though. A notable exception is the use of “fake news” by American newspapers during the Second World War to expose the propaganda of Joseph Goebbels, presented by the Nazi hierarch as innocent report news (Lepore 2018: 454). It is only in the early 2000s that the expression underwent a deep process of re-interpretation that it is still running nowadays (Gelfert 2021).

In the first decade of the 21st century, “fake news” became increasingly connected to satire. TV shows like *The Colbert Report*, broadcasted from 2004 until 2015 through the American cable TV channel Comedy Central, were increasingly referred to as fake news

show. As Michael Gettings (2007) says, this kind of TV shows satirizes traditional news media «by reporting in a style similar to network and cable TV news, but it amplifies their biases, mistakes, and deficiencies to ensure that viewers hear them loud and clear» (pp. 26-27). It is only in 2016, for the aforementioned reasons, that "fake news" acquires the Collins Dictionary's meaning, so that regaining more the initial sense of the 19th century. But as many reporters well know, there was another twist in the history of fake news.

This new addition is due to the unscrupulous use that Donald Trump made of the term after his electoral triumph. In fact, the American tycoon started using "fake news" as a way to referring to media outlets – such as the *New York Times*, *Washington Post*, and CNN – that he just did not like (Silverman 2017)[1]. In time, fake news slowly became a weaponised expression that aims at delegitimizing newspapers, acknowledged institutions and epistemic authorities in general (more on this further on).

Even though the history of the term had many turning points, there has been an outbreak of works in journalism and academia that make use of a fixed meaning of "fake news" in order to describe a phenomenon that deserves to be studied in its own right. Those studies include works from political science, law, communication, informatics all the way to philosophy. The common idea is that fake news diagnoses a sort of threat for our informational environment that needs to be addressed. Scholars do often disagree on the novelty aspect of this threat. Concerning philosophy, for instance, some (like Pepp et al. 2019a) emphasize how the infrastructure of social media is something essential to fake news, thus giving to the phenomenon a distinctive nature compared to the kind of problematic news that might have been spread on traditional media. On the other side, others (like Mukerji 2018) suggest not to confuse an empirical link – namely, the fact that recently fake news stories were communicated almost exclusively through the internet and especially social media websites – with a conceptual connection: there may be nothing in the concept of fake news that entails a close and necessary connection with social media infrastructure.

[1] Actually, Trump claims he invented the term "fake news". He stated this during a press conference in October 2, 2019, when he said: «I call the fake news now corrupt news because fake isn't tough enough. And I'm the one that came up with the term – I'm very proud of it, but I think I'm gonna switch it to corrupt news» (Beaujon 2019). Even though Trump did not strictly invent the term, it is true in a sense that the last twist in its significance – the weaponized sense of the term – is due to his contribution to some extent.

But apart from sociological contentions like this, the vast majority of philosophical endeavours shares the same intent: finding a satisfactory definition of fake news. Recently, the conceptual analysis of fake news has in fact drawn a great deal of attention from several philosophers. Especially in the past lustrum, many epistemological endeavours have been carried on in order to find an answer to the *vexata quaestio*: what is fake news?

The aim of this chapter is to summarize the current epistemological debate on fake news and to highlight some major problems of the prevailing positions. In §2 I present some unescapable challenges that trouble any definitional project of fake news: meaning instability, conceptual necessity, and weaponization. Here you find also the response that some scholars have tried to give to those same challenges. In this same context I anticipate my position to give a more precise idea of the scope of my project. Then, in §3, I briefly summarise the two main positions in the debate in epistemology on fake news: the conservative view and the privative view. After that, I dedicate a specific discussion to each of the two positions, underlining their main characteristics and the criticisms to which they are exposed. Criticisms include both counterexamples to conceptual analysis of fake news and critiques of conceptual necessity (more on this soon). This latter part takes up most of the chapter, and goes from §4 to §9. Finally, in §10, I recap the chapter's findings.

2. The definitional project and its challenges: Habgood-Coote's thesis

The underlying idea of conceptual analysis is that any philosophical question around a certain concept, like fake news, presupposes a clarification of that same concept. Before addressing issues around the potential menace of fake news for democracy or the possible antidote to it, we need first of all to determine what fake news actually is. More clearly, conceptual analysis for fake news aims at establishing the necessary and sufficient condition for a news item to count as a fake news item. We know that we possess a good bundle of conditions if it resonates well with our intuitions about particular cases. When our conditions correctly determines, according to our intuitions, whether or not something falls under the concept x in a fair number of cases, we might

reasonably think that our definition of x is good[2]. To search for necessary conditions means to stipulate a notion of fake news that is not excessively narrow: it does not exclude from the concept too many probable tokens of that concept. It is the other way round for sufficient conditions, which are fair enough if our concept is not overinclusive: we do not want to include cases that are straightforwardly outside the scope of the concept in question.

Now, when conceptual analysis turns to fake news it may stumble on obstacles that are related to the very notion under discussion. Regarding this point, Habgood-Coote (2019) summed up three main challenges of defining fake news that should prompt us to abandon the term, according to him. Firstly, Habgood-Coote argues that fake news faces a problem of meaning instability: it has no fixed descriptive meaning due to its historical variability. There is also a plethora of definitions on the academic market that seems to point to disagreement even among experts. If we also consider that ordinary speakers too have conflicting belief about the meaning of fake news, Habgood-Coote continues, we reach a dead end: both ordinary and expert intuitions are out of use. In addition, the tricky historical evolution of the term makes impossible to trace it back to its original use.

Secondly, Habgood-Coote raises an issue of conceptual necessity. As he notes also elsewhere (Habgood-Coote 2022), we can rise the following questions concerning fake news: do we really need a term like this in our philosophical vocabulary? Does it add any useful descriptive resources to our language? Or do we already have other concepts – like disinformation, misinformation, propaganda and others – that sufficiently supply our intellectual needs? Habgood-Coote's answer to the latter question is affirmative. Whatever the meaning of fake news, it remains a superfluous concept, and we just do not need it. Usually, scholars involved in the epistemological debate have mainly addressed the challenge of meaning instability, giving a conceptual analysis of the notion of fake news in terms of necessary and sufficient conditions. Regardless of problems of extension adequacy (that do exist), these accounts have often taken for granted that the question of conceptual necessity was already settled. That is, that we do

[2] There is no unanimous consensus on which insights matter most. Someone claims that expert opinion takes the lion's share (Putnam 1975), while others value more intuitions of ordinary speakers of a linguistic community (Burge 1979). Usually the best definitions are the ones which meet the intuitions of both groups.

need fake news as a theoretically fruitful concept. So far, though, this claim is disputable because we already possess plenty of concepts to talk about many pathologies of communication fake news is meant to refer to, Habgood-Coote concludes.

Lastly, there is the question of political weaponization. Habgood-Coote, and other authors as well (Talisse 2018, Coady 2021), argues that fake news has assumed the role of an epistemic slur with propagandistic purposes: you can deem something as "fake news" in order to discredit its source and to rally the troops against it. Fake news has turned into an accusation of media manipulation to be affixed to one's political counterpart. The outcome of such an involution is that our use of the term in every conversation, though backed by good intentions, may end up reinforcing its slurring content. According to Habgood-Coote, the very existence of a meaning instability is the first reason that allows this perverse usage of fake news: propagandists may target the term at whoever they want, given that there is no consensus on who (or which) should be the target of fake news. The second reason is that, despite lacking a stable descriptive meaning, fake news has often a clear evaluative content. Usually, fake news points to something inherently bad, unreliable or biased.

Some experimental evidence corroborates Habgood-Coote's stance. As a part of study on top-of-mind associations, van der Linden, Panagopoulos and Roozenbeek (2020) asked people with a political affiliation, from a national sample of the US population, to report the very first thing that comes to their mind when they think of the term "fake news"[3]. The three authors hypothesized that, in the mind of partisans, "fake news" might be associated with ideologically discordant media sources. They were right: the result of their study shows evidence that both liberals and conservatives associate traditionally right-wing and left-wing media sources, respectively, with the term "fake news". In short, people tend to associate "fake news" to media outlets that cover the opposite part of their political spectrum. Therefore, the combination of lack of stable descriptive meaning and presence of a strong, negative evaluative content enables the propagandistic usage of fake news. These three challenges over the term "fake

[3] Top-of-mind associations amount to mental representations of an issue, and specifically they are the first words, thoughts, or symbols that come to the subject's mind when they are prompted with a cue (Clarke et al. 2015; Nelson et al. 2000). In the case of van der Linden et al. (2020), the cue is the term "fake news".

news" – meaning instability, conceptual necessity, weaponization – should prompt us to advocate its abandonment, Habgood-Coote concludes.

However, some scholars involved in the debate on fake news (Brown 2019, Pepp et al. 2019) disagree with Habgood-Cote's reprimand. Let us see some of those discordant opinions.

2.1 Vs meaning instability

Contra Habgood-Cote, Brown (2019) claims that philosophers, as well as journalists, do not need to refrain from using "fake news", provided that they use it *cum grano salis*, and rejects Habgood-Coote's challenges one by one.

Concerning meaning instability, the fact that "fake news" has a disputed descriptive content makes this term similar to many others – such as "neoliberal", "manipulative", "social deprivation" etc. – that we nonetheless retain in our vocabulary, despite the lack of consensus over their meaning. Lack of stable descriptive meaning is not in itself a sufficient reason to discard none of those terms, "fake news" include. Furthermore, even though the adjustment of the descriptive content due to expert disagreement may be a slow process, this does mean it is infinite: eventually, we will come up with a satisfactory definition of the concept. In the meantime, Brown continues, philosophers and journalists may keep talking about fake news with the aim of raising questions concerning moral and epistemic values around, say, our interactions with online information. What Brown intends to get across is that, whether or not it has a fixed meaning, a concept may be fruitful for our intellectual discussions in many different ways. So when we use fake news with this goal in mind, there is no reason to care about meaning instability: we can simply use it to scrutinize our moral and epistemic practices.

Pepp, Michaelson and Sterken (2019) agree with Brown in stating that Habgood-Coote's arguments are not decisive to advocate for the abandonment of the term "fake news". If the metasemantic conclusions of Habgood-Coote on fake news are valid, then we would be forced to drop many other terms that describe complex and social phenomena, like "combatant", "spousal abuse", "intellectual property", and "insider trading". And this hardly qualifies as a desirable scenario. Nonetheless, Pepp, Michaelson and Sterken claim that it may well be that there are some use of fake news

that should be definitely discarded. Regarding this point, the authors bring the two examples of fake news being used to talk about satire and comical news and the propagandistic usage denounced by Habgood-Coote.

2.2 Vs conceptual necessity

Scholars disagree with Habgood-Coote's objection of conceptual necessity too. According to Brown, it is questionable that fake news is conceptually superfluous. Brown puts forward two arguments to support fake news conceptual necessity. The former simply denies that there is no social phenomenon that fake news can describe. In fact, terms like "propaganda", "lie", "bullshit" and others do not capture what seems to be a specific kind of deception brought about fake news: the mimicry of traditional news stories. As we shall see, some authors – the so called "privativists" – believe that a fundamental property of fake news is that it imitates some formal properties of genuine, traditional news. If that is the case, we may now have good reasons to retain fake news in our philosophical vocabulary, given that no other concept enables us to talk about this kind of deception.

The latter argument of Brown, instead, goes as follows. Even if fake news did amount to something for which we already have a concept, there would be value in even understanding how this concept interacts with existing ones. For instance, if we end up accepting that the work made by the concept of fake news is already performed by the concept of propaganda, it would be interesting to understand in what kind of relation the two concepts stand. Is it fake news a particular case of propaganda? If so, how does it differ from ordinary cases of propaganda? And so on and so forth (Brown adds further arguments against the challenge of conceptual necessity. For instance, he wonders whether Habgood-Coote is in the right epistemic position to judge that fake news is an unnecessary concept. For the sake of simplicity, though, I shall avoid those other arguments here).

Concerning the challenge of conceptual necessity, the bulk of Pepp, Michaelson and Sterken's argument is to use fake news' negative evaluative content strategically. Assuming that Habgood-Coote is right – i.e. fake news has only a fixed evaluative content – we may tip the scale in our favour by striving to restrict the class of objects that deserves that negative evaluation as much as possible. Thus, "fake news" would

have at least one reason to be helpful (and so conceptually necessary). This kind of operation, the authors continue, is actually the converse of projects to reclaim social slur terms. In reclaiming a slur someone, usually members of the group targeted by slurs, takes a term with a clear referent and gets rid of its negative evaluative content (see for instance Brontsema 2004 on the reclamation of "queer"). In reclaiming fake news in such a way, instead, we do the opposite: we choose a term equipped with a strong evaluative content and provide it with a restricted descriptive content so that we can target only those (fake) stories that actually merit the negative evaluation. The practicality of this operation rests on two assumptions. Firstly, the negative evaluative content associated with "fake news" may remain stable for a while, and since people are unlikely to accept dropping the term altogether then reducing the scope of the term may be wiser in the short term. Secondly, the restriction of the descriptive content may be carried out in limited contexts, like academia or journalistic production, that seem to be more receptive to this kind of worries, and thus more successful in performing the whole operation (think about experimentation of inclusive language in the scientific community and press).

2.3 Vs weaponization

Regarding Habgood-Coote's last objection, Brown eases the concern about weaponization of "fake news". According to him, to think that literally every use of "fake news" in any conversation conjures up the ideology of media manipulation is a claim that should be supported by empirical evidence. As he notes, «it is far from evident that scholars who publish scientific journal articles about fake news, for instance, have a major impact (if any) on the public's perception of major news outlets» (Brown 2019: 151). To assert this, Habgood-Coote needs sufficient empirical evidence of the alleged impact of scientific production on ordinary speakers' top-of-mind associations with "fake news"; which is currently not the case[4]. Besides, the risk of

[4] Following Brown, I believe that it is far more likely that negative associations with "fake news" are consistently formed in the specific context of politics. There are at least two reasons to think so. The first one is mere correlation. The explosion in popularity for the term "fake news" and its subsequent weaponization coincided with the 2016 US Presidential campaign: one of the most talked-about political election of the third millennium. Secondly, we have plenty of evidence that negative associations around fake news are very common in politically-laden conversations among partisans. Political affiliations is in fact one of the most influential factors in determining people's beliefs about which news sources are

being weaponized seems to be very natural for every concept that somehow leans into politics. Why do we not abandon them all even though they face the same risk of "fake news"? The reason of why we retain them, Brown argues, is that we think that «such a risk is outweighed by the political gains we make by continuing to use such concepts» (ibidem). But then we may apply the same logic to "fake news" as long as its usage can provide us with important conceptual gains, such as democratic evaluation of the ways in which governments or Internet companies should regulate social media and discussions about moral responsibilities of social media users.

Instead, Pepp, Michaelson and Sterken agree that "fake news" is frequently used in propaganda, as denounced by Habgood-Coote. But recognising this fact, once again, is not a sufficient reason to do without the concept in question. The reason is that we normally feel confident in using many other concepts that already ended up being exploited as a form of communicative and political weapon. Among these ones we may include some like "propaganda" itself, "biased", "lying media" and so on. The fact that politically polarized groups, for instance, use "propaganda" as a tool for political accusation does not *ipso facto* drive us to discard the term in every context: the acknowledgment of the first kind of occurrence does not necessarily exclude the peaceful coexistence with the other ones. In addition, Pepp, Michaelson and Sterken criticizes Habgood-Coote's idea according to which lack of precise descriptive content is the main driver behind the derogatory use of fake news. Propagandists have an easy time accusing news sources of being fake news not because the target of this accusation is often vague and indefinite, but due to the prominent role played by the evaluative content of the term. The reason is that disdain towards the target, and not facts concerning it, is what counts the most for propagandists and their flattering audiences.

2.4 Taking stock: how shall we think of fake news?

In light of this debate, there are a few conclusions we can draw. The first is that Habgood-Coote's worry over conceptual necessity of fake news seems warranted. Many definitions of fake news fail to provide necessary concepts, because they do not add much more to the theoretical resources we already have. Furthermore, there is a whole set of issues with the inclusion of conditions concerning mental state of fake news

"fake" (van der Linden et al 2020; Tong et al. 2020; Michael & Breaux 2021), even though is not the only one (see Axt et al 2020 for psychological drivers behind fake news accusation, for instance).

producers in fake news definition. As we shall see in the next paragraph, the majority of fake news definitions comes from hybrid account and suffer from the problem highlighted by Habgood-Coote: they run the risk of being unnecessary duplications of other, more perspicuous notions. This is true despite Brown's observation according to which fake news, for some authors, coincides with a new kind of intellectual imposture that consists in the imitation of traditional news stories. In a further reply to the received objections, Habgood-Coote (2020) points out that the sense and extension of the mimicking action of fake news is not fully established by proponents of intellectual imposture model. In contrast, there exist an array of different interpretations for the kind of imitation at play, each of which would identify a different notion of fake news (more on this in §9, in this chapter). But even assuming that a firm position can be found on the notion of mimicry, problems still exist with this thesis, which suffers from counterexamples targeting requirements for imitation. As we shall see, it may be rational to dub as fake news also news items which blatantly deviates from ordinary journalistic writing and look (more on this in chapter 2 and 3).

However, what is missing in Habgood-Coote's analysis is the explanation of why this happens: why does fake news ends up overlapping with other notion that already describe pathologies of information? One aim of this work is to answer this question and to encourage a definition of fake news that can solve the problem of conceptual necessity. I maintain that the main reason of why fake news struggles to be an independent concept is due to a hidden assumption that underpins many positions in the debate. The assumption is that news is just information of some kind. According to this point of view, a news item is a bundle of information that we get from news sources, and whenever that piece of information is somehow ruined we actually have fake news. Out of metaphor, the ruin occurs when information is false or potentially misleading. Other times news is corrupted because it results from epistemically flawed attitudes on the part of the news producer. More often than not they are deceptive attitudes of some sort, even though there is no plain consensus on what the main deceptive attitude is (Grundmann 2020). Beyond the nuances of the different accounts, the model remains often the same: news is a kind of information, and so fake news is a kind of corrupted or deviated information.

One of the fundamental thesis of the present work is that news is not the same as information. News and information are two related yet independent entities. I argue that information is an epistemic entity, while news in my view is a *sociological* entity. The former has to do only with meaning and truth values, and for someone it has a structural connection with the concept of knowledge (Floridi 2010, 2011). News is an entity of a different kind. It is partly related to meaning and truth values, since it incorporates information, but it does not reduce solely to them, nor it has a structural connection with knowledge. Furthermore, it seems to have some relations with a bunch of fragmented dictionary's properties, like recency and virality – even though those properties alone do not explain why news and information come apart. The main reason for their difference is that the possibility of being of news hinges on the prior existence of a specific institution, whereas information is not in this dependent relationship. The institution that enables news to exist is that of journalism, which in turn sets out the norms (or rules) that news should observe in order to be dubbed as such. Fake news is therefore news that is significantly at odds with those norms. It is not that they always reflect the appearance of real news, since fake news can present outward features that manifestly deviates from the style of journalistic writing. Rather, as readers we have flexible expectations of journalistic writing: we can tolerate even a certain degree of deviation from the norms of journalism, as long as it is done for a just cause. But what is the purpose of this whole proposal for fake news?

Following Chalmers (2020), the present work involves a project of homonymous conceptual re-engineering: it aims at establishing a new meaning for an old expression ("fake news"). It is a stipulative project because it does not want to capture what we ordinarily mean by "fake news", but is instead trying to disentangle this notion from concepts with which it might be confused. It is a kind of operation that departs from conceptual analysis. The latter tries to capture a concept we actually have and use, while stipulation is more like a revisionary analysis: it is the project of setting out a notion of fake news that is useful, or that we need; that is, an ameliorated concept of fake news (Haslanger 2000/2012, 2020)[5]. The scope of the stipulation is to endow the

[5] This project does not even have to accept that there is any unified concept of fake news. In fact, I believe as Pepp, Michaelson and Sterken do, that introducing and spreading the meaning of a term takes some time for linguistic communities, especially when the term in question undergoes many changes in its usage in just a few years (possibly three changes from early 2000s to 2016). Therefore, I do not regard

philosophical vocabulary with a concept that is useful and sufficiently distinct from others. Hence, what the present work is trying to do is something like a redesign of fake news for academic usage only. As in the case of reverse reclamation proposed by Pepp, Michaelson and Sterken (above, p. 15), here we restrict the descriptive content of fake news on purpose, but with the goal of isolating only its real contribution to the philosophical debate. This should help securing fake news from unwanted contaminations with similar concepts. Besides, by relegating fake news to the context of academia as much as possible we may discourage its weaponized use, for the university world seems to be more receptive to the threat carried by epistemic slurs, as Pepp, Michaelson and Sterken already noticed.

The project of re-engineering fake news has the advantage of not being too expensive for ordinary speakers. Coming up with a new term is always less palatable for speakers who may find it difficult to implement a neologism in their conversations. Instead, homonymous conceptual engineering takes advantage of the fact that «old words are entrenched and often have associated prestige» (Chalmers 2020: 11). On the other hand, this kind of conceptual engineering has the downside of having to demonstrate why the proposed meaning is the legitimate one. There may be multiple proposals of stipulation for fake news coming from different scholars. Why should the one in this dissertation be the legitimate one? Concerning this point, my project of stipulation fits into the framework of the so-called conservative view of fake news, even if paradoxically it is inspired by the works that fall under the privative view (more on both labels in the next section). Thus, the stipulative definition that will be presented in the third chapter is not something completely new, but is intended to be a sort of refinement of one of the positions in the debate. Since it is compatible with some of the definition proposals already in the field, the project of conceptual engineering performed here does not present itself as the only legitimate one, but it is based on insights already provided by the debate.

The next sections will briefly present the two main positions in the epistemological debate on fake news. Sections §4, §5 and §6 are devoted exclusively to the conservative view and the possible problems it faces. Since it represents the predominant position in

---

meaning instability as a problem *per se* nor do I set out to provide a conclusive solution to this challenge here.

the debate I will devote much more space to it. Section §7, §8 and §9 are instead dedicated to the other faction, that of the authors who support the privative view.

3. The current debate: conservative vs privative accounts

As I have stressed, that there is no consensus on the exact meaning of fake news. However, roughly two factions can be identified within the current philosophical debate over the definition of fake news: the conservative view and the privative view[6]. The basic idea that distinguishes the two factions concerns whether or not fake news belongs to the news genre.

According to the conservative view, fake news really is a species of news, albeit of a flawed kind. The vast majority of the accounts under this label are hybrid ones. This denomination stems from the fact that many scholars sharing the conservative approach set out two different conditions for a news item to be a fake news item. Those conditions has to do with the two dimensions of truth and truthfulness, so that fake news is news that lacks both truth and truthfulness. News lacks truth if it is either literally false or implicitly conveys falsehoods. News lacks truthfulness if its producer either has the intention to deceive the audience in some way or is indifferent to the truth of the news content. Some accounts in the literature, though, derive from the hybrid ones by subtraction of one of the two conditions. Therefore, these derivative accounts only require that a news item lacks either truth or truthfulness in order to be considered as fake news. In any case, even these positions share the idea that fake news is a certain type of news, and therefore belong to the overarching conservative approach. Here is the general schema for conservative definitions of fake news:

CD: Fake news is news that lacks truth and/or truthfulness

As we shall see, the conservative view is more vulnerable to the criticism of conceptual necessity promoted by Habgood-Coote. In addition, some of the conservative

[6] The latter denomination was coined by Bernecker, Flowerree and Grundmann in the introduction to the collection *The Epistemology of Fake News* (2021), while the former is a terminological invention of mine. They instead contrast the latter with the so-called hybrid view, but as I am going to say right here hybrid accounts are only the prevailing version of the conservative, but they do not represent the totality of its members.

definitions of fake news either fall prey of counterexamples or present controversial aspects that go beyond the issue of conceptual necessity.

The privative view of fake news is built upon an opposite principle. For these authors fake news does not belong to the same genre of news, but is a story that only masquerades as news. The denomination of this position comes from the fact that these accounts insist on the privative function of the adjective "fake" in fake news. Privative adjectives are modifiers that usually deny that the object designated by the whole noun phrase belongs to the noun extension (more on this in §5). If so, fake news do not designate a genuine news item, but only a forgery that does not have the right history of production (which for news usually coincides with having been produced according to standard journalistic practices). Forgery lies in the fact that, despite not being authentic, fake news is able to reproduce the appearance of genuine news. For many privativists, imitation of genuine news is the real act of deception carried out by fake news: it misleads its recipients about its status as news. Lack of truth is typical but not necessary for fake news, according to the privative view. Since fake news is a story that was not produced following good practices it is much more prone to inaccuracy and falsehood. Here is the general schema for privative definitions of fake news:

PD: Fake news is information (or a story) that is deceptively presented as news but it is not

The privative view of fake news has gained increasing acceptance in the last few years, but since it is a relatively young position there is not much criticism of it in the literature. In a further reply to his commentators, Habgood-Coote (2020) has emphasised the lack of agreement among privativists on the type of imitation performed by fake news. In addition to this observation, I will supplement it with a personal reflection on the limits of the concept of imitation within privative view. This concept is linked to that of counterfeit (or forgery): fake news is counterfeit/forged news precisely because it imitates the appearance of real news, according to privativists. However, privativists' idea of imitation is vulnerable to counterexamples due to the tricky nature of the concept of forgery/counterfeit. As we shall see, privativists do not make sufficiently clear the requirements for a piece of news to be a credible imitation of genuine news.

In the next two paragraph I will single out conservative and privative accounts for a more in-depth analysis of the two positions. For both views I present some known arguments against their definitions and supplement them with novel arguments when required (this will mostly occur for criticisms of privative accounts). The critical arguments will consider both problems related to meaning instability (first of Habgood-Coote's challenges) and those that hinge on conceptual necessity (second of Habgood-Coote's challenges).

4. The conservative view of fake news

Proponents of conservative view regard fake news as a subset of news itself. For the authors belonging to this group, news is usually propositional content – asserted or implied – that represents one or more events occurring in the world. This content can manifest through various mediums such as articles, images, videos, and any means by which to express propositional content[7].

From the conservative perspective, fake news constitutes asserted or implied propositional content, albeit of a problematic nature. For instance, a prominent fake story during the 2016 US election asserted that Pope Francis endorsed Donald Trump for President (Silverman 2016). However, it's essential to note that producers of other news items may conceal problematic content implicitly. For example, they might suggest a causal link between two unrelated events, thereby prompting readers to derive a false implicature. Jaster & Lanius (2018) and Croce & Piazza (2021) make an example that sounds like the following one. Consider a newspaper headline like this:

(1) After the arrival of the migrants, the crime rate in town has increased by 10%.

Suppose, then, that (1) is published in a context in which we know for sure that the crime rate was raising due to local factors and even before the migrants'arrival. This is a case in which the falsehood of a piece of news is clearly located on implicit level.

Most of conservative accounts are also hybrid accounts (or hybrid views). The label "hybrid" refers to the double condition required to a piece of news in order to be identified as an instance of fake news. Normally, hybrid accounts provide for an

[7] Belonging to the conservative group only requires believing that fake news is a kind of news. This does not necessarily imply believing that news is itself only asserted or implied propositional content – as conservative authors normally do.

objective condition – a requirement that applies to the propositional content of news – and a subjective condition – which instead applies to the disposition and/or the actual intentions of the news producer. A classic candidate for the objective condition is a requirement that concerns the falsehood or the misleading nature at some level of news content, while the subjective one usually calls into question some epistemically bad or psychologically defective attitudes on the part of the producer, as well as her hidden goals.

This twofold characterization, combined with the focus on propositional content presented as news, gives us the answer of hybrid view to the first challenge of meaning instability: fake news is a kind of news – which in turn is propositional content – that is flawed in two respects, namely in its content and in the way it is disseminated. According to this perspective, fake news is news deprived of something, namely of true or straightforward content and a proper attitude in disseminating it. A representative case is the definition put forth by Jaster and Lanius (2018):

> [Fake news] lacks truth and truthfulness. More specifically, fake news is either false or misleading (lack of truth) and it is propagated with either the intention to deceive or an utter disregard for the truth (lack of truthfulness) (Jaster and Lanius 2018: 208).

The first feature of fake news in the definition of Jaster and Lanius is the lack of truth – the objective condition – which affects both explicit and implicit level of meaning, separately or both together, thus accounting for trickier cases like (1). The other feature consists in the lack of truthfulness – the subjective condition. Almost all hybrid definitions of fake news state that producers must possess some kind of deceptive intention, but its nature varies across authors. Jaster and Lanius, as well as other scholars (Rini 2017, see below; McIntyre 2018), hold that the deceptive intention may concern the news content. In this case, falsehood and/or misleadingness of the content is exactly what is deceptively got across. Other authors, instead, maintain that producers of fake news intend to deceive consumers about the source of a piece of news, and not about its content. This is the case of Fallis and Mathiesen's proposal (2019), according to which there must be the intention to deceive the audience into believing that a

reported story was a genuine one; that is, a story that has gone through the standard modern journalistic process[8].

Another possible subjective condition is the one advanced by the second member of the disjunction in the subjective condition of Jaster and Lanius's definition. Indeed, lack of truthfulness includes either the intention to deceive the audience or the more complex attitude of bullshitting: cases in which a piece of news is transmitted without any concern for the truth at all. According to the adapted interpretation of Frankfurt (2005), a bullshitting producer does not want to necessarily deceive about news content; instead, she wants to deceive readers about her attitude towards the truth, pretending to be sincerely concerned about it when, in reality, she is totally indifferent. Even if the Frankfurtian account is the most prominent one in the literature on fake news, in another work Jaster and Lanius (2021) specify that they intend bullshit in a different sense, according to which «to bullshit is to contribute an utterance as an answer to a question under discussion without concern that what one says or communicates is based on evidence» (p. 22).

Whatever the interpretation under discussion, the appeal to the bullshitting attitude should enable hybrid definitions to label some specific cases as fake news too. Among these there is clickbaiting. The most cited example of this kind are the fabricated news produced by the Macedonian teenagers before the US presidential election in 2016. These teenagers claimed that they were just exploiting the partisanship of Trump supporters with the sole intent to make money by means of advertising; they were completely indifferent towards the truth of news content and they did not care about deceiving news consumers (Silverman & Alexander 2016). Therefore, Jaster and Lanius consider the false or misleading products of Macedonian farms as a case of bullshitting and, consequently, of fake news too.

Another hybrid definition, slightly different from the first one, is the one offered by Regina Rini:

> A fake news story is one that purports to describe events in the real world, typically by mimicking the conventions of traditional media reportage, yet is known by its creators to be

[8] Strictly speaking, the account of Fallis and Mathiesen (2019) belongs to the privative view of fake news. But these typological labels are often just an approximation and it is much more convenient to mention this feature of Fallis and Mathiesen's position here for the sake of the argument. We shall return on their proposal in §7.

> significantly false, and is transmitted with the two goals of being widely re-transmitted and of deceiving at least some of its audience (Rini 2017: E-45).

Leaving aside the mimicking feature for a moment, here we have again an objective condition which requires that the content of the story is significantly false and a subjective one, in the form of two goals: 1) achieving a wide retransmission and 2) deceiving at least some consumers. Apart from the deceiving intention, which we have already discussed, the first goal is a recurrent one in the domain of hybrid definitions of fake news. It says that in order to be labelled "fake news", a news item, among other things, must be intended to reach a minimum number of people. Call this the minimum audience condition for fake news.

The primary reason to incorporate a certain version of the minimum audience condition in the definition of fake news has to do with the notion of news itself. For many conservative authors (Rini 2017; Gelfert 2018; Mukerji 2018, Jaster and Lanius 2018, 2021), it is somewhat implicit for news that it is read or at least accessible to a vast audience (more on this in the next chapter). Since an oral statement made in private to one person cannot be news, they say, a fortiori it cannot be fake news. But the incorporation of the minimum audience condition may be driven by other reasons as well. For instance, if someone believes that widespread dissemination of false beliefs is an essential aspect of fake news, then they may include some version of the minimum audience condition in their characterisation of fake news, since news items which reach a vast audience are more likely to induce more falsehoods than ones which have lower circulation[9].

Rini is not alone in advocating the minimum audience condition. Gelfert (2018), for instance, shares Rini's remark on the constrain of widespread circulation but adds that «more is required for an item of disinformation to count as fake news; namely, a certain

[9] Anderau (2021) rejects this argument because it assumes that a greater quantity of people harmed by fake news will always trump the quality of how badly they are harmed. He then provides an example of a fake news item read by a single guy who gets so enraged by the news content that he goes wild. This counterexample should demonstrate that «fake news can cause great harm even if it reaches only a single reader, which makes it questionable why one would set up a threshold of minimum readers» (Anderau: 208). I do not think that Anderau's argument is a knockdown one. For it is sufficient for a conservative account to claim that we are measuring the extent of damage caused by fake news only in terms of false beliefs, and not in terms of consequence that may (or may not) result from those beliefs at a later stage. Therefore, the wider the audience the greater the dissemination of false belief, regardless of what happens next to that same audience.

measure of success in *realizing* the goal of widespread circulation and uptake» (p. 102). Gelfert's point is that the sole intention of reaching a wide audience is not sufficient *per se* to convey fake news, since the attempt in question must be somehow effective. Recalling an episode that dates back to his time as a physics student, Gelfert tells of a quirky man who tried to disseminate, with a blackmail, his semi-professionally produced pamphlet denying Einstein's relativity theory. Gelfert suggests that, even if the man was successful in manufacturing his pseudo-scientific text by mimicking the look of an official scientific paper, the attempt would have been hardly counted as an instance of fake news, since «the manner of delivery undercut any chance of achieving wider circulation» (Gelfert 2018: 102). Regarding the same question, Croce and Piazza (2020) partially develop Gelfert's analysis by saying that an assertion can be regarded as the assertion of fake news «provided that it is made over a medium that is at least conducive to the attainment of that goal [i.e. the widespread circulation]» (p. 51). In summary, for the hybrid view the satisfaction of the minimum audience condition should be either successful or at least prospective, when fake news content is spread through media (e.g. social media) that virtually allow producers to meet that goal.

However, conservative accounts are not all hybrids. There are also conservative proposals that are formed by subtraction of one of the two conditions of the hybrid view. We can call them "derivative accounts" of the conservative view. Mukerji's definition of fake news is an exemplary one:

> Fake news is [Frankfurtian] bullshit asserted in the form of a news publication (Mukerji 2018: 929).

Here there is no objective condition: fake news does not need to be false or misleading in its content. Indeed, Mukerji does not consider the falsity or the misleadingness of content neither a necessary condition nor a sufficient one, because, he claims, even a fabricated story that accidentally comes to be true appears to us to be fake, while honest mistakes committed by journalists do not strike us in the same way: good faith is incompatible with a bullshitting attitude, which in turn is the sole subjective requirement of Mukerji's definition.

We can situate Gelfert's account in the group of 'monoconditional' proposals too. Gelfert (2018) proposes to understand fake news as in (G_1), but refines it in Gelfert (2021) as in (G_2):

> (G_1) Fake news is the deliberate presentation of (typically) false or misleading claims as news, where the claims are misleading by design (Gelfert 2018: 108).

> (G_2) Fake news is the deliberate presentation of manipulative and misleading content as news, where the content is manipulative and misleading by design (Gelfert 2021: 320).

Gelfert reverses the operation carried out by Mukerji: both of his definitions are deprived of a strong subjective condition, while retaining an objective requirement. There is no intention to deceive/bullshit the audience. Instead there is something weaker: the «deliberate presentation» of content despite its *intrinsic aptitude* to manipulate or mislead the consumer. In other words, on the part of producers there is the expectation that the content in question is likely to fool the consumers in some way, but nevertheless there is no intention to carry out an act of deception. That is why Gelfert states that fake news content is manipulative and misleading «by design». The type of manipulation he refers to is that which prompts the reader of a fake news item to a reaction of some kind. It does not need to be the formation of some sort of belief, but it can mean pushing people into the arms of a political party or inciting hatred against social groups of some kind.

Also Grundmann (2020) has put forward a definition of fake news that is deceptive-intention free:

> News that is produced or selected in general ways such that it has the robust disposition to lead, at the time of publication and under normal conditions, to a significant amount of false beliefs in a significant number of the addressed consumers (Grundmann 2020: 8).

This account defines fake news in terms of the harmful effects on beliefs that it typically has on news consumers. What is at stake here is not really if fake news is actually misleading for its consumers, but only if it has the disposition to do so in the vast majority of cases; that is, misleading under epistemically normal circumstances. Those circumstances include characteristics of the audience (in terms of smartness,

competence, cognitive bias, attention span etc.) and medium of message transmission. For instance, a highly knowledgeable person is more likely to realise that a news item is fake than someone who is unfamiliar with the information world and its dynamics, all other things being equal.

As long as what is required for a piece of news to be fake is its disposition to affect consumers, not the intentions or other mental states of its producers, the account of Grundmann is non-hybrid. Still, it is a conservative account. Grundmann is rather explicit on this point, for he argues against privative accounts which maintains that fake news is counterfeit news on the ground that "fake" is a privative adjective. As far as I know, Grundmann is the first author who take a stand against the general interpretation of fake as a privative adjective in the context of fake news. Rather, he defends a conception of fake news as a kind of genuine news, where "fake" indicates only that news is somehow deceptive (more on this in chapter 3).

In the following section I critically discuss the hybrid view by arguing that we should abandon any reference to subjective conditions, because it is either unnecessary (in the sense of Habgood-Cote 2019) or conducive to unpalatable outcomes for our conceptual endeavour.

5. Doing without subjective conditions

When we are embarked in conceptual analysis we usually strive to offer definitions that are neither overinclusive nor overexclusive: they must not include cases that should be excluded from the extension of the concept and vice versa. Normally, a good way to test a definition is looking for one or more counterexamples that, intuitively, seem to be unduly included or excluded by the same definition. That is one of the main strategies that I undertake in this section, which criticizes the subjective conditions put forth so far by conservative accounts (hybrids and derivatives). This section gathers and systematizes counterexamples against requirements on intentions in the literature, as well as providing new arguments down the road. This first round up of arguments deals with the challenge of meaning instability. That is, the present section criticizes (one of) the conditions that the conservative view adopts to determine the descriptive content of fake news.

According to the overview provided in §4, broadly speaking, there are two kinds of intention for the subjective condition of conservative view, the former of which can be further divided in two:

a) The intention to deceive the audience, which can be targeted to the news content or the news source so that we actually have two distinct kinds of deception;

b) The intention to bullshit the audience[10].

As already outlined by Grundmann (2020), if we combine hybrid definitions that select only one of the two former intentions together with the disjunctive proposal of Jaster and Lanius, we obtain four different possibilities:

A1. The intention to deceive the audience about news content (Rini 2017);

A2. The intention to deceive the audience about news source (Fallis and Mathiesen 2019);

B. The intention to bullshit the audience (Mukerji 2018)

C. Intention (A1-2) or intention (B) (Jaster and Lanius 2018).

I will single out and discuss each of them in the next subsections. Afterwards, I say a few words about the minimum audience condition.

5.1 Deceiving the audience about content (A1) or source (A2)

To understand why we do not need the intention to deceive the audience (in both senses), perhaps it is useful to grasp the reason for why we may need it. At first glance, in fact, this intention should enable us to exclude specific cases of false/misleading news from the domain of fake news; namely, the ones in which an undesirable outcome is not simultaneously backed by malevolent attitudes, just like in the case of accidental journalistic mistakes. According to Jaster and Lanius (2018), for instance, even if reputable media sometimes do issue false or misleading statements, they usually do so inadvertently, because they normally aim at the mark of truth. Fake news producers, instead, aim at falsity and so they deliver false or misleading statements on purpose:

[10] Although there exist different interpretations of bullshitting, here I limit the discussion to the Frankfurtian account. See Fallis and Stokke (2017) for another interesting interpretation of bullshitting related to contributions to the question under discussion within a conversation.

lack of truthfulness set fake news apart from journalistic errors. I find this claim unconvincing for several reasons.

First of all, it seems to me that what drives us to the exclusion of such cases is not properly the good faith of the producer in general. Rather, it is the social identity of the producer in question – namely, a journalist – what persuades us in doing so. We cannot infer the same conclusion without the former assumption according to which a journalist is, by definition, someone who normally aims at the truth. In other words, it seems that the supposed function of the subjective condition under scrutiny is actually performed by a former judgement on journalism as social enterprise which remains behind the curtain in hybrid proposals[11]. If so, the intention to deceive the audience is a subordinated condition, and therefore is useless and may potentially lead us astray, since it covers the hidden assumption about social identity of producers.

Secondly, stated in this way, the condition may end up being too narrow: it does not include suspect cases of false information in the domain of fake news. Consider a modified version of the Pizzagate story: let us call it PIZZAGATE* as in Mukerji (2018)[12]:

> (1) Pizzagate*
> A white supremacist devoted to the cause of truth claims she has discovered an obscenity perpetrated by Hillary Clinton and other Democrats. She said she has broken the secret code behind Clinton emails which would reveal the existence of a paedophiliac ring headed by Clinton and her fellow party members. Obviously, none of this is true. Still, the white supremacist decides to publish her 'findings' online, without any intention to spread false information, and still believing that she is just rendering a service to the truth.

If the subjective condition about deceiving intentions is at play, we cannot deem PIZZAGATE* as fake news: the false story was presented without the intention to deceive the audience. However, the conclusion reached so far appears to us to be completely counterintuitive, since we may have a strong inclination to label the claims of the white suprematist in PIZZAGATE* as fake news too. I should note that exclusion of such examples may not be a critical problem for a conceptual endeavor; however, hybrid

[11] A judgement like this is made unfold, instead, in some alternatives to the hybrid view. See for instance the account of Pepp, Michaelson and Sterken (2019) mentioned below, and my own proposal in chapter 3.

[12] This example is also present in Pepp et al. (2019).

view does not really do enough to justify an outcome that, at first sight, seems at least controversial.

This last remark drives us to a further concern. The intuition that PIZZAGATE* must be considered as fake news despite its lacking malevolent intentions seems to suggest that, in applying the label "fake news", we are more sensitive to *process and outcome* of news production rather than mental state of producers. After all, if one of the main concerns around the spread of fake news is the danger of being systematically deprived of true or justified beliefs, we may not be really interested in producers' intentions, provided that we end up forming true or justified beliefs anyway as a result of exposure to shared news content[13]. At the end of the day, what really matters is maintaining a positive ratio between knowledge and false/misleading information, regardless of the intention of those who participated in the process of beliefs' generation. Obviously, the whole thing depends on the nature and the scope of the debate: the same conclusion does not automatically apply for debates within which conceptual endeavour is deeply concerned with speaker's mental states, such as the definition of lying.

I should stress that the objections made here equally apply to both A1 and A2, since the main factor under discussion is the intention to deceive broadly conceived, while the difference between A1 and A2 hinges on objective condition (content vs source). However, I will elaborate more on the intention to deceive on the source in §8.2, since the account of Fallis and Mathiesen deserves a dedicated section.

5.2 Bullshitting the audience (B)

In my view, there are also many motivations for not having a bullshitting intention within the definition of fake news. First of all, I am not really convinced by how the Frankfurtian interpretation of bullshitting should enable us to account for clickbaiting. Recall that, according to the hybrid view, cases of clickbaiting like Macedonian farms are instances of bullshitting attitude inasmuch as the clickbaiter does not care about the truth of what she is disseminating. What the clickbaiter really cares about is making

[13] About ending up with justified beliefs, we may think about a case in which a consumer, say Robin, follows a rule of thumb that states "always believe in what newspapers report". Whenever Robin should inadvertently run into a piece of news which, unbeknownst to her, was spread by a local newspaper with the intent to deceive the audience; nevertheless we may end up saying that Robin's subsequent belief is warranted anyway, since in ordinary cases we are justified in trusting newspapers. Obviously, we are assuming, among other things, that Robin's epistemic environment is an 'healthy' one overall (e.g. Robin must not live under a repressive regime that prevents freedom of the press).

money by means of advertising, and therefore she publishes only news that is apt to fulfill this goal, regardless of their truth.

Furthermore, Jaster and Lanius (2021) add that Macedonian teenagers «may well want to deceive their audience about their attitude towards the truth. In fact, deceiving their audience in this way seems to be a necessary means for getting clicks» (p. 22). So far, then, the hybrid view deeply connects the achievement of clickbaiters' goal – that is, generating as many clicks as possible – to the fulfillment of the bullshitting attitude, the latter being also a necessary condition for the former.

Now, it seems to me that the attitude of indifference towards the truth is neither a necessary condition nor a sufficient one for achieving clickbaiters' goal. It is not a necessary condition because we can imagine without contradiction all sort of successful websites, in terms of clicks, which are at the same time not indifferent to the truth of the published content. In addition, the producer's lack of concern for the truth does not entail that she is also unconcerned with whether her audience believes what she says. Instead, clickbaiters *must care a lot* about sparking interest into the audience because

> if the audience don't find the claims even remotely plausible, then why on earth would they be clicking on them under the guise of them being news? The point is that clickbait that wants to generate internet traffic by offering misleading claims as news must be in the business of also wanting to actually mislead its clientele, since that is precisely what is going to ensure that this strategy is successful (Pritchard 2021: 49).

The indifference towards the truth on its own, then, is not even a sufficient condition for getting clicks. To be potentially successful, the clickbaiter must also try to be credible or at least attractive for the news consumer. Hence, she will try to mislead her audience, even though this may not be her only intention in play or the overarching one.[14]

Secondly, the Frankfurtian account of bullshitting is not able to deal with PIZZAGATE* in a satisfactory way. Mukerji treats this case as the example of an unwitting bullshit. Following an insight of Ladyman (2013) about pseudoscientists, Mukerji claims that «the attitude of indifference need not be a conscious one. Rather, it

[14] Notice that I am only criticizing the hybrid view with regards to its strategy to include clickbaiting under the fake news umbrella. I am not by any means suggesting, between the lines, a new fake news definition that should include the misleading intention in order to account for clickbaiting. I will offer my own definition only in chapter 3.

can be located at a deeper level of the agent's psychology» (Mukerji 2018: 941). A producer may have a first-order self-representation as someone truthful and at the same time be not sufficiently sensitive to evidence in her inquiry as someone truthful. Then, if she deceives herself into believing in her distorted first-order representation, we can conclude that she is bullshitting unwittingly, as Mukerji says. Consequently, we may treat the sincere author of PIZZAGATE* story as one that has been fooled by her own self-representation, thus ending up being a bullshitter without her knowledge. However, as Pepp et al. (2019) have already noticed, the author's self-deception is not enough to label PIZZAGATE* as a case of bullshitting, since this attitude also requires the bullshitter to try to come across as someone who is concerned with the truth; that is, the bullshitter also wants to deceive her target into believing in her interest towards the truth. If so, Mukerji ought to conclude that the author of PIZZAGATE* intends to unwittingly cover up the fact that she is bullshitting unwittingly[15]. Now, this is a really intricate intention, and sounds deeply unnatural to maintain for a subject. We may prefer solutions to PIZZAGATE* that do not force us to postulate such a sophisticated intention.

Thirdly, the intention to bullshit does not work for a further variation of Pizzagate story, PIZZAGATE**:

> (2) Pizzagate**
> A fear monger white supremacist claims she has discovered an obscenity perpetrated by Hillary Clinton and other Democrats. She said she has broken the secret code behind Clinton emails which would reveal the existence of a paedophiliac ring headed by Clinton and her fellow party members. Obviously, none of this is true. Still, the white supremacist decides to publish her 'findings' online, even though she knows that her claims are false, for the sheer sake of spreading disinformation.

Again, PIZZAGATE** seems to deserve the label of fake news, even if it does not fall within the scope of Frankfurtian bullshits at first glance. Instead, it is a clear-cut case of lying. Mukerji admits that PIZZAGATE** is more difficult to reconcile with his account, but then goes on to say that this objection can easily be countered because lying and bullshit are not exclusive categories. As Frankfurt explains, to bullshit is to make

[15] Perhaps Mukerji implied that the target of the covering-up intention of bullshitting is the producer herself, but that is far from a way out. If so, the producer in question would be somehow bullshitting herself instead of the news consumer, and then the notion of unwitting bullshit would become useless for our purposes since it does not guarantee a relation between a producer and a consumer.

assertions "without paying attention to anything except what it suits one to say" (Frankfurt 2005: 60). A bullshitter may happen to know that a given assertion is false but choose not to pay attention to that fact since he is indifferent to the truth. In that case, he is not only a bullshitter. Under the standard definition of lying, Mukerji argues, he is also a liar because he asserts something he knows to be false. Accordingly, PIZZAGATE** can be construed as a case of Frankfurtian bullshit, too.

But asserting something the speaker knows to be false is not a sufficient condition of lying: it is just a necessary one. The standard definition of lying, to which Mukerji refers to, is composed of three conditions. Specifically, A lies to B if and only if there is a proposition *p* such that: 1) A says that *p* to B; 2) A believes that *p* is false; 3) by saying that *p* to B, A intends to deceive B into believing that *p* (Stokke 2013, Mahon 2015). Mukerji's solution contemplates (1) and (2) but leaves out (3), therefore he cannot conclude that bullshit and lying match each other completely. Besides, Mukerji's remark on the first lines of his paper rules out his own solution. Indeed, he says that «unlike a liar, who tries to convince us that a given statement he believes to be false is true, a bullshitter simply does not care whether what he says is true» (Mukerji 2018: 921).

5.3 Disjunctive accounts (C)

Definitions like the one presented by Jaster & Lanius (2018) cover the entire spectrum of cases, because here fake news results either from an intention to deceive consumers or from bullshitting attitudes. Their definition, then, faces both objections made in §2.1 or the ones presented in §2.2. Moreover, this disjunctive account may raise a problem of theoretical unity. In fact, the "lack of truthfulness" condition risks conflating distinct speaker attitudes that are flawed for different reasons; and this is usually a problem for a definition (McKeown-Green 2009).

First of all, two senses of truthfulness must be distinguished. In truth-conditional semantics, truthfulness relates to whether a statement corresponds to reality or is factually accurate. It is concerned with the relationship between language and the external world and indicates an 'objective' feature of the statement. On the other side, in philosophy of language truthfulness is sometimes framed in terms of a speaker's commitment to honest communication: a person is truthful if they communicate only what they believe to be true, even if their belief turns out to be factually incorrect. When

Jaster and Lanius write about (lack of) truthfulness are clearly referring to the latter meaning of the term: truthfulness as a commitment to always say what one believes to be true. Then, lack of truthfulness is just the absence of this commitment.

This second sense of truthfulness, though, overlaps significantly (and dangerously) with the notion of sincerity in speech act theory. In speech act theory, the felicity of many speech acts depends on the speaker's intention to communicate sincerely. A speaker is sincere when they express their genuine beliefs, feelings, or intentions through their statements. We have a case of insincerity, then, when there is a discrepancy or an incompatibility between the mental state expressed by the utterance and the actual mental state of the utterer (for other conceptions of insincerity, see Pepp 2018). So although sincerity has a broader application to all speech acts – and includes emotional and intentional alignment, not just factual accuracy – when it comes to beliefs and truth-oriented speech acts, such as assertions, truthfulness and sincerity overlap. If a speaker says, "I think it will rain tomorrow" and genuinely believes this, they are both (pragmatically) truthful and sincere. As a consequence, in the specific context of beliefs, untruthful communication and insincere communication indicate the same thing: the speaker's utterance does not reflect their genuine inner state.

If so, when Jester and Lanius include both an intention to deceive and a bullshitting attitude in the "lack of truthfulness" condition, they may confuse insincere communication with deceiving intentions. For instance, I am insincere when I assert "It is raining outside", thereby expressing the belief that it is raining outside, but at the same time I believe that it is not the case. Bullshit requires a kind of insincerity too, albeit of a slightly different kind: I believe neither that what I have said is true nor that it is false, nonetheless me saying that *p* expresses my belief that *p*. But both cases have nothing to do with a further, independent intention to deceive my addressee. I may be bullshitting without intending to deceive about the content of my utterances, and vice versa. Therefore talking about lack of truthfulness (aka sincerity) with regard to deceiving intentions seems inappropriate – unless the deception concerns one's own mental states, but the examples brought by Jaster and Lanius clearly focus on the content of the utterance. This brings me to the concern of theoretical unity: the "lack of truthfulness condition" groups together different phenomena under the same label (see also Grundman 2020 for a similar objection).

5.4 The minimum audience condition

In §1 I mentioned the goal of reaching a wide audience as one possible subjective condition for fake news. I then renamed this goal "the minimum audience condition" (MAC henceforth). There are several problems with existing formulation of MAC for fake news. Here we will take an in-depth look at some of them.

First of all, some formulations of MAC suffer from counterexamples. Specifically, the most vulnerable are those that claim that fake news is published with the aim or intention of being widely retransmitted. Rini (2017) is one of those accounts. Here is a counterexample for this kind of formulations:

> (3) The chain
> Imagine that Caroline, a computer expert, wants to dupe her friend James. Caroline crafts a news item through artificial intelligence about a fictitious political scandal. Then, she shares the preview of the article as an Instagram story that is visible to her followers. The preview includes the headline, an AI-generated image and a brief caption. Just before sharing it, though, she changes the privacy settings of her Instagram story so that James is the only follower of her profile that is allowed to see that story. Caroline does this because she wants to reach and deceive James and him alone. She does not intend to reach or deceive anyone else with her fake preview. Afterwards, James sees Caroline's story but does not fall for its content. Indeed, unbeknownst to Caroline, he sends a screenshot of her story to his friend Michael with the goal of convincing him, and him alone, of the truth of the political scandal mentioned in the fake preview. Michael in turn sends the same screenshot to another friend of his with the same goal James and Caroline had – convincing just the immediate recipient – and the chain keeps stretching until it reaches, say, fifty people.

Even though in THE CHAIN the scandal is false news transmitted with deceptive intentions to a lot of people, each person on the chain sent it with the goal of reaching only their immediate recipient. Thus, Rini is compelled to dismiss this scenario as not a case of fake news, even though this judgment clearly sounds wrong.

The second problem is one of indeterminacy. Whatever the formulation at stake, the threshold of how big the audience should be to satisfy MAC cannot be determined once and for all. What is the minimum number of readers of a fake news item that fulfils the condition? On this point, privative accounts such as the one of Pepp, Michaelson and Sterken (2019) makes it clear that the required number seems to vary in accordance

with the news item's field of interest. For instance, if a news item involves events concerning local politics in a small town, then the people who are likely to care about it might be just the inhabitants of that town, say a few thousand people. Therefore, the threshold for MAC will be a more or less large subset of that population (say one hundred inhabitants). By contrast, if the news item in question deals with a story that encompasses national or even international events, then the threshold will be far higher. The problem of threshold indeterminacy – the impossibility to fix an exact number of people for MAC – remains therefore unsolved.

The third problem of the minimum audience condition is a bit more complex. We said in §4 that one of the main reason to accept a minimum audience condition for fake news is to account for the harm of dissemination of false beliefs. I claim that each formulation of the minimum audience condition falls short of accounting for this hazard of fake news. To begin with, prevailing formulations of MAC do not consider the divergence in roles among four different categories of agents in the news cycle: producers, consumers, receivers and propagators (Croce and Piazza 2019). Producers are those who arrange the news item and make it public for the first time. Consumers are the news item's 'passive' addressees: they read it but do not necessarily believe it (they may disbelieve it, ignore it and so on), nor do they undertake any further interaction with it. Receivers are consumers who end up believing the news item. Finally, propagators are those who share the news item through social media, thus contributing to its spread – without necessarily believing it themselves (thus resembling consumers rather than receivers). One consequence of this subdivision of roles is that a news item may reach an extensive audience of consumers without causing them to believe it – that is, without encountering receivers. Consequently, reaching a wide audience of pure consumers may not be enough to render fake news harmful[16].

To solve this problem, we may simply tweak the minimum audience condition and say that a news item must be designed and spread in a way that is conducive to the goal of reaching a wide audience of receivers. However, even this second version of MAC is not fully satisfactory, since I will argue that there exist a group of cases in which a news

[16] Someone might object that fake-news producers may not always wish to produce a belief in the consumer as their intended effect. Perhaps they are aiming at other, non-doxastic states (Mercier 2020, Bergamaschi Ganapini 2021, Hannon 2021). But if we endorse an account of fake news as something that need only lead to purely non-doxastic forms of acceptance, then we should also reject the idea that fake news is dangerous because it spreads false beliefs.

item is not designed to reach a wide audience of receivers and yet still strikes us as an instance of fake news. I call these cases outlandish fake news. They are stories that, beyond simply being untrue, are highly unlikely to be true given our general knowledge about the agents and facts at stake. Outlandish fake news is deeply at odds with the way in which the world normally works, so to speak. Some examples of outlandish fake news are Pope Francis' alleged support for Donald Trump (Ritchie 2016), the supposed plan of the British Ministry of Defence to requisition Wembley Stadium to cook the world's biggest lasagne (Waterson 2020), and the fabricated statements of Keanu Reeves saying that Hollywood elites use the blood of babies to get high (Mikkelson 2017). All those examples may have reached a vast audience of consumers, but I find it hard to believe that there are as many readers who have bought into it (i.e. receivers). How could anyone fall for such tall tales? Their manifest absurdity is inevitably unconducive to the satisfaction of our revised MAC because the average consumer will refuse to believe them. Therefore, outlandish fake news poses a problem for accounts which incorporates MAC.

In conclusion, I claim that the minimum audience condition, as it has been designed so far, is not compatible with the dangerous nature of fake news. The reason is that it either fails to distinguish news believers from mere news consumers (who are not harmed) or excludes a particular subset of dangerous news items: outlandish fake news.

6. The problem of redundancy

In the previous section we have seen in detail the problems of conservative definitions with regard to the first definitional challenge: the challenge of meaning instability. Now we turn to the problem of conceptual necessity, i.e., seeing. "fake news" as an unnecessary concept (Habgood-Coote 2019): we already have plenty of concepts for talking about deceit, miscommunication and other phenomena that the notion of fake news should deal with. Besides, those concepts may fare better than "fake news" since they have clearer and stabler meanings in ordinary language and they are much less easily weaponized. The present section does not add too much to Habgood-Cote's argument, but it purports to make it clear why his worry is well-founded when it comes to the conservative view of fake news.

To understand why the conservative view would put forth a redundant concept, maybe an analogy is helpful. Suppose someone claimed that we need to talk about "mosaic lies", where that means lies made up of fragments of truth pieced together to create a misleading picture. While we might acknowledge that this picks out a particular phenomenon, we would probably find it redundant. While "mosaic lies" is evocative, it does not expand on what other, established terms already convey. For instance, we can just talk about half-truths. So, we do not need the concept of mosaic lies, because it does not pick out a distinctive – important, worth focusing on – kind of lies. The main problem with fake news understood in a conservative way is pretty much the same of mosaic lies: we already possess theoretical resources that account for phenomena that it should deal with without bringing up the notion of fake news. We do not need fake news to serve in this role.

Concerning instances of false information deliberately spread to mislead consumers, for instance, we normally possess a term that serve the purposes of accounting for those cases, and it is "disinformation" (Fallis 2015, 2016). Furthermore, the conceptual contribution of disinformation is far more definite and unambiguous than that of fake news, since the former benefits from a debate devoid of weaponization and political contaminations. To differentiate their notion of fake news from the one of disinformation, defenders of the conservative view have two alternatives: either sheltering their model by dropping the subjective condition on deceptive intentions – as in the case of derivative accounts – or turning to bullshitting-like attitudes. As we shall see immediately, both solutions are problematic.

The former is far from being satisfactory. In fact, it raises again the same criticism of conceptual necessity: the common term "misinformation" already exemplifies cases of dissemination of false content, regardless of intent to mislead (Fallis 2016). We do not need "fake news" as a conceptual duplicate of misinformation. At first glance, the latter solution of turning to bullshitting-like attitudes seems more promising, and deserves unpacking.

Normally, a notion such as disinformation (or misinformation) can account for a phenomenon of epistemic deprivation (Croce and Piazza 2019). Epistemic deprivation means that consumers who are systematically exposed to false information run the risk to increase their total share of false or unjustified beliefs, while simultaneously

decreasing their knowledge. Such a scenario potentially undermines the self-determination of individuals even within a democratic regime, since the accumulation of false information may push them towards harmful choices in sensitive areas, like personal health or economic status. Epistemic deprivation is a threat to democracy that notions such as dis/misinformation are able to capture.

When fake news ends being a duplicate of dis/misinformation it enables us to account for the same phenomenon of epistemic deprivation. However, fake news is not always conceived as a duplicate of dis/misinformation within conservative accounts. Some of them, like that of Mukerji (2018), involve the producer of the fake news acting as a Frankfurtian bullshitter, regardless of the truth value of the news item in question. But then, uncontrolled dissemination of fake news understood in this way may give place to a phenomenon different from epistemic deprivation. The reason is that bullshit is not necessarily false. Its hallmark lies in the attitude of indifference towards the truth of what the bullshitter affirms or publishes. Hence, epistemic deprivation cannot be the real threat for an informational environment plagued by people who are indifferent to the truth of their claims. In this case, the real menace is a general loss of interest for the truth and its authority in human speech, as Frankfurt claims. It is not the risk of collecting false beliefs. Rather, it is like decreasing the general commitment to knowledge transmission in human interaction – a no less significant damage than epistemic deprivation.

Accordingly, if fake news enables us to identify the risk of the general loss of interest for truth-conducive statements in online communication, then defenders of conservative view have another string to their bow. Since other rival concepts of fake news – included disinformation and misinformation – cannot account for such potential threat to democracy, we may conclude that the notion of fake news elaborated by conservative accounts does not designate a redundant concept after all. If so, it might be argued that, contrary to mosaic lies, Mukerji's definition does pick out a distinctive kind of bullshit; namely, the one which has the form of a news publication. But it is just a sleight of hand. Actually, by doing so we are surreptitiously using the notion of fake news as a vicarious concept for bullshit. That is, we did not attribute a new conceptual meaning to fake news; we have just named a particular instance of bullshit "fake news". If so Mukerji's definition of fake news is not conceptually necessary in our sense, since it is

the notion of bullshit that does all the work here: the danger of a lack of interest for the truth is fully and uniquely explained by the Frankfurtian account of bullshit. Having the form of a news publication does not add anything, per se, to the dimension of harm of bullshit[17]. In other words, fake news here might at most be a subcategory of bullshit, but it cannot constitute a new conceptual category on its own. Once again, we do not need fake news to serve in this role.

Here ends the discussion of the problems for the conservative view. So far we have a bunch of controversial points concerning the subjective condition for fake news, as well as an issue of conceptual necessity: conservative interpretations end up being superfluous, because the meaning of fake news they support overlaps with the one of other, less controversial concepts. In the rest of the chapter I am going to present the rival faction on the debate, which is the privative view of fake news. The plan for the next sections is the same. First a general overview of privative accounts, then a critical discussion.

7. The privative view of fake news

The second main faction in the fake news debate is the privative view. Privative authors maintain that a fake news item is a piece of information or a story that has the appearance of news but it is not. Privativists locate the 'fakeness' of fake news in its status as news rather than in its truth value or in the producer's flawed intentions. For the majority of these authors (Fallis & Mathiesen 2019; Lazer et al. 2018; Pepp et al. 2019), fake news is fake because it lacks the required pedigree of being produced by standard journalistic processes. For someone else (Pritchard 2021), the fakeness of fake news lies in the tendency of its source to convey inaccurate information; in contrast to genuine news sources, whose aim is delivering correct information to the audience. In both cases, though, privativists agree that fake news does not necessarily presents falsehoods in its content. Nonetheless, fake news is *typically* false, because non-compliance with journalistic standards leads more easily to falseness.

Both positions, though, agree that fake news is not really news, and ground this assumption on the interpretation of "fake" in fake news as a privative adjective.

[17] This does not rule out the possibility that this further characterization adds something important about bullshit in a different respect. For instance, it may say something about the notion of assertion, bullshit or news or the relationship among these three things.

Privative adjectives are noun modifiers that usually allow for the following inference: "If X is Adj(N) than X is not N". Thus, "fake news" designates an object that does not really belong to the extension of noun "news", according to privativists. That is why these authors often contrast fake news to genuine news, the latter class being news that does not lack either the pedigree of being produced by standard journalistic processes or the tendency to convey accurate information. Fallis and Mathiesen (2019) explicitly employ this distinction in their definition of fake news:

> Our preferred definition is that fake news is counterfeit news. In other words, a story is fake news if and only if it is not genuine news, but is presented as genuine news, with the intention and propensity to deceive (Fallis and Mathiesen 2019: 8).

And then they go on saying that

> By genuine news, we mean stories that have gone through the standard modern journalistic process [...]. That is, genuine news has been produced by professionally trained reporters, fact checkers, and editors, who are attempting to provide fair and accurate accounts of current events (ibidem).

These lines of Fallis and Mathiesen (2019) raise many points. Their definition is given by way of contrast: fake news is news that is not genuine. Plus, there is also a subjective condition, since fake news is presented as news with an intention to make people believe that it is genuine news when it is not. Fallis and Mathiesen (2019) compare the counterfeit nature of fake news to that of coins. Just like a counterfeit coin is a coin deprived of the requisite history of production of authentic coins, and is so forged with the intention to deceive, fake news is counterfeit news that is made to fool people into thinking that it has the same history of production of genuine news. This is something different from intending to deceive about the content of the news. Fake news may not be intentionally deceptive about its content, but it always is on its status as news, according to Fallis and Mathiesen.

Not having the requisite history of production is the main fault of fake news also according to other privative authors. According to the account of Pepp, Michaelson and Sterken (2019), fake news can be defined in the following way:

> Fake news is the broad spread of stories treated by those who spread them as having been produced by standard journalistic practices, but that have not in fact been produced by such practices (Pepp et al. 2019: 69).

Again, not having been produced by standard journalistic practices (i.e. not having right history of production) seems to be the main requirement for a news item to be fake.

The first difference with the account of Fallis and Mathiesen is that here there is no intention to deceive whatsoever. Pepp, Michaelson and Sterken do not build any subjective condition into their definition mainly because it would exclude cases like PIZZAGATE* from the domain of fake news (see above, p. 24). They characterize journalism as a dynamic social institution and pinpoints a set of ideals that every journalistic venture ought to achieve. Instances of such ideals are the commitment to the interest of readers, a scrupulous verification of information, transparency about sources and conflicts of interest, and so on. So fake news is a story that has not been produced in accordance with those ideals. Regarding this point, Pepp, Michaelson and Sterken add two clarifications.

Firstly, to qualify as produced by standard journalistic practices, a process of story production need not maximally fulfill all the ideals mentioned before. In order for stories to possess a stronger claim to being genuine news, it is sufficient that their story-production processes live up to these ideals to a significant extent; otherwise, when such processes are substantially out of line with these ideals, they are unlikely to be deemed as genuine news. However, due to the inherent vagueness that pertains to the subject matter, the authors do not advance sharper guidelines that might enable us to carry out this distinction.

Secondly, accomplishing (at least) a good deal of such ideals does not ensure the trustworthiness of a piece of news; that is, following the standard journalistic practices is not a sufficient condition for ensuring that no false or misleading news content is ever spread. This was somehow already included in the former clarification, insofar as Pepp, Michaelson, and Sterken did not say that news with a proper story-production processes are *ipso facto* genuine news; rather, they hold that such news has a stronger claim to being genuine news. Sometimes unfavourable factors get in the way of journalism and are simply inevitable.

Not every privative account bases the contrast between fake news and genuine news on the observance of journalistic standards, though. Pritchard (2021) assumes that news is a kind of information that one gets from a news source. A news source, in turn, is any source of information that is designed to convey accurate information to others about recent events[18]. From this basis, Pritchard makes a distinction between news sources that are epistemically good and those that are epistemically bad. The former are news sources that are especially good at doing their job because they follow standard journalistic practices; it is other way round for epistemically bad news sources. But here is where Pritchard's account drifts apart from other privative ones. Even though epistemically bad news sources do not follow journalistic practices in their doing they still remain news sources, albeit of a low quality. The reason is that definition of news sources here is not motivation-oriented. News sources is designed to convey accurate information in the sense that it aims at doing so. News sources which do not adopt good practices, like journalistic ones, are simply doomed to many more failures in their outcomes than sources which instead adopt journalistic standards. But they are not necessarily fake news sources.

Despite their denomination, fake news source are not really news sources. For Pritchard, fake news is not a kind of news because its source does not meet the requirement of being aimed at conveying accurate information. In contrast, it is designed to spread misinformation. Again, the crux of the matter lies in the intent. Whether or not fake news manages to misinform the audience does not count. What counts is that fake news sources are always in the business of misinforming news consumers; and this very design is what keeps them separated from epistemically bad news sources. On this point, Pritchard directly calls into question previous privative accounts which confuse fake news with genuine news that is epistemically problematic, according to him. The reason is that Fallis and Mathiesen (2019) and Pepp et al. (2019) both incorporate good epistemic procedures into the notion of news itself , but this is «to conflate news with a good epistemic pedigree with news *simpliciter*, and thereby to

[18] Pritchard clarifies that news sources of interest to us must be variable in terms of the kind of news delivered. This addition rules out sources of information that cannot deliver more than one message (like the post-it left on the fridge). Moreover, we are considering only news sources that are responsive enough to the world in terms of information delivered. The traffic light gives us accurate information but it does so by determining it, and not by responding to an external stimulus (as in the case of news reporting).

fail to capture the fact that there can be news that has a poor epistemic pedigree but is no less news as a result» (Pritchard 2021: 55).

Regardless of this point, all three accounts share the idea that fake news mimics the appearance of genuine news. The mimicking action is one of the main acquisition of the privative view. Even though other accounts attributable to the conservative view might have maintained the same point incidentally, privativists have made it almost predominant. The reproduction of the aspect of genuine news is an essential aspect of fake news, according to them. But what exactly is being imitated by fake news?

In this regard, there is no plain consensus among privative accounts. Fallis and Mathiesen do not commit themselves to any specific position on the matter. They simply recognise more cases of imitation by quoting a line from Tandoc et al. (2018: 147) that says that fake news appropriates the look and feel of genuine news, including how websites look, how articles are written and how photos include attributions[19]. On the same topic, Pepp, Michaelson and Sterken observe that people identify journalistic writing by means of outward features — such as the use of headlines, bylines, captions, certain prose styles and diction — and other kinds of looks that are customary for both offline and online print, such as Twitter feeds. That is, according to them there exists a set of outward features that is a hallmark of journalistic writing, and it is precisely what comes to be reproduced by fake news:

> Hence, by employing these looks and superficial features, content originators may produce stories that look like genuine journalism, whether or not they are, and whether or not content originators are aware that they are making their stories look like genuine journalism (Pepp et al. 2019: 77).

The mimicking action of those producers may not be conducted intentionally (that is why Pepp, Michaelson and Sterken's account does not include any deceiving intention in their definition of fake news).

Other privative accounts, such as Lazer et al. (2018) and Anderau (2021), mainly talk about the imitation of the form of genuine news publication. For these authors, it is

[19] Someone may notice that Rini's definition too incorporates the imitation of genuine news appearance (above, p. 24). However this trait of fake news is presented as Rini only as typical, and not necessary. Therefore imitation of genuine news is not a core feature of fake news, according to her definition.

important that fake news is published, broadcasted or shared in such a way as to look like real news. Specifically, Anderau (2021) emphasizes how fake news imitated format can change over time due to the ongoing evolution of media. Even though news format is dynamic, Anderau argues, what matters is if the fake news item could be seen as real news at the time and place it is published. This means that fake news does not need to assume the appearance of traditional news media: the format of social media or from other communication channels are suitable, provided that are consumed by an audience that considers them legitimate formats for real news.

Finally, Pritchard (2021) is elusive on the kind of imitation at play for fake news. On one hand, he asserts that fake news is something that masquerades as news in order to spread misinformation (he does not detail what this masking consists of). At the same time, he maintains that an ineffective attempt to reproduce the format of real news may still qualify as fake news, provided that it is designed to convey inaccurate information. Unlike Fallis and Mathiesen (2019), then, a fake news item may not have a propensity to deceive about its status as news. Pritchard gives us the example of a lousy fake news website that is very easily spotted by the average news consumer. Even though the overall activity of this website is ineffective – because it fails to misinform its audience due to its inherent sloppiness – that does not entail that its product is not fake news. This stance implies that a purposely misleading piece of news that does not imitate well the appearance of real news may be still dubbed as fake news in Pritchard's account.

Up to now, not many criticisms of the privative view have been developed in the literature. The reason is twofold. First of all, the privative view is a relatively recent position in the debate and has not yet been subjected to an extended critical scrutiny. Secondly, many authors have shifted their attention from the definitional project to other urgent issues surrounding the topic of fake news. In the rest of this chapter I put forward some arguments against the privative interpretation of "fake" in "fake news" and I reinforce the critique of conceptual necessity proposed by Habgood-Coote (2020). Although my position on fake news will completely unfold only in chapter 2 and 3, some of the arguments developed in the following sections anticipate my main thesis.

8. The imitation of genuine news: a critique

In this section, I make some observations on the privative assumption that fake news reproduces the appearance of real news. Firstly, I point out that we do not have enough evidence to say that this is actually the case. Indeed, there is also evidence to the contrary: fake news can adopt a style of its own, which deviates from that of real news. Secondly, I point out that this privative assumption is vulnerable to counterexamples.

8.1 What does empirical evidence say?

So far, the overarching majority of privativists endorses the idea that fake news mimics the appearance of real news. This act of imitation is instrumental in making the story or the piece of information in question look like genuine news when it is not. Privativists further claim that lack of truth is not a necessary feature of fake news, but only a typical one, since forged news is made without following the correct journalistic practices, and therefore is much more exposed to the risk of being false. If the last claim is correct, we may run a sort of backward-looking search: to test the hypothesis that fake news imitates genuine news, we may check if false stories presented as news do replicate the appearance of real news in a sufficiently large number of cases. If so, the privative assumption about mimicking activity of fake news would have an empirical confirmation. However, it seems that there is currently insufficient evidence either to confirm or deny this assumption. To explain this point, something must first be said about the search for detection methods for fake news.

The predominant focus of research on the language of fake news lies within the realm of natural language processing (NPL), primarily geared towards developing tools for the automated discernment between genuine and fake news (Rubin et al. 2015; Shu et al. 2017; Asr & Taboada 2018; Bondielli & Marcelloni 2019; Oshikawa et al. 2020; Zhou & Zafarani 2020). This literature often defines fake news as news that is false. The immediate consequence of adopting such definition is that collecting dataset of fake news that serve to test machine learning systems is relatively easy. It is sufficient to draw on the archives of fact-checking organisations, like PolitiFact or Snopes, and media outlets that have an editorial staff dedicated to fact-checking (like BuzzFeed). But there are also disadvantages as well. The first and foremost is that the veracity-based definition leaves out cases of fake news that are literally true but deceptive under

some other respect. There are many case of news whose content is true but it ends up misleading readers, or news that manipulates consumer through other ways (e.g. by inciting racial hatred). Furthermore, veracity approaches usually merge cases of news that is unintentionally false with those of news that is deliberately false. Therefore, the backward-looking search mentioned before either does not cover all case, because false news is only a subset of fake news, or does not sufficiently diversify cases of false news. Once this clarification has been made, it is useful to see the conclusions of some studies using a veracity-based approach.

Through a study of three data sets and features that capture the style and the language of articles, Horne and Adali (2017) found that fake news appearance shares more characteristics with satire than real news. After defining fake news as stories that are known to be false and are from well-known websites that intentionally spread misinformation, Horne and Adali computed many content-based features on each data set and categorized them into three categories: stylistic, complexity, and psychological[20]. Below is a list of the most significant results of their study:

- The content of fake and real news articles is substantially different. Real news articles are significantly longer than fake news articles and fake news articles use fewer technical words, smaller words, fewer punctuation, fewer quotes, and more lexical redundancy.
- Fake news articles seem to be filled with less substantial information demonstrated by having a high amount of redundancy, more adverbs, fewer nouns, fewer analytic words, and fewer quotes.
- When looking at just the titles of fake and real news articles, there is an even stronger dissimilarity between the two: fake news titles are longer than real news titles and contain simpler words in both length and technicality. Fake titles also used more all capitalized words, significantly more proper nouns, but fewer nouns overall, and fewer stopwords (i.e. words that are automatically filtered out by natural language processing systems due to their insignificance).

[20] Roughly, stylistic features amounts to differences in text, syntax and grammatical elements among article content and headline; complexity features measure the overall level of intricacy of articles and titles, in terms of both sentence structure and word readability; finally, psychological features are connected to sentiment analysis of the magnitude of negative or positive emotions elicited by words in each text.

- Fake content is more closely related to satire than to real news content[21]. Specifically, both satire and fake news use smaller, fewer technical, and fewer analytic words, as well as, fewer quotes, fewer punctuation, more adverbs, and fewer nouns than real articles. Further, fake and satire use significantly more lexical redundancy than real articles.

The study from Horne and Adali shows that fake and real articles are usually distinguishable, specifically in the title of the articles, thus questioning the mimicking feature of fake news and apparently arguing for the opposite thesis. The study of Horne and Adali is not the only one coming up with certain results. Carrasco-Farré (2022) and Grieve and Woodfield (2023) also shows that fake news, as well as other pathologies of communication, has their own peculiar features in terms of style, cognitive effort required to read it, appeal to emotion and so on.

However, the results of these studies do not settle the issue once and for all. The definition of fake news proffered by Horne and Adali, as well as the one from Grieve and Woodfield, does include cases of intentionally false news, therefore it does not face the second limitation of veracity-based approaches (above, p. 38). However, the first limitation applies here: the definition of Horne and Adali neglects cases of misleading and manipulative fake news, and the same goes for Grieve and Woodfield. It is not that they do not recognise their existence, they simply do not deal with those cases. So proponents of privative views may suggest that their study does not rule out the possibility that fake news that is only misleading does imitate the appearance of genuine news. At the same time, it is true that other empirical investigations seem to point to the opposite direction. After conducting a stylometric inquiry on hyperpartisan news, Potthast et al. (2017) conclude, for instance, that «the style of fake news has much more in common with that of real news than either of the two has with satire» (p. 238). These results might be due to differences in data sets, type of evaluated features, size of the samples and starting assumptions about definition of news categories. Therefore the two studies cannot be compared so easily. But the point is that, at least in theory, we have two opposite conclusions on the same matter.

[21] Horne and Adali define satire news as «stories that are from news sources that explicitly state they are satirical and do not intentionally spread misinformation» (Horne & Adali 2017: 760).

Nonetheless, the main takeaway here is that the mimicking feature of fake news is far from being empirically settled. At best, we find ourselves with contradictory data that do not enable us to pass judgment on the matter; at least until representatives of the privative view, who bear the burden of proof, do not prove otherwise.

8.2 A forgery is (not) always a forgery

In §7 we saw that there is disagreement among privativists over the degree of success that a fake news item has to obtain in reproducing the appearance of genuine news. Fallis and Mathiesen (2019) say that a piece of fake news must have the propensity to deceive about its status. Fake news which does not really look like real news may not be fake after all. In contrast, authors like Pritchard (2021) maintain the opposite. Lousy fake news is still fake news, because a forgery remains a forgery regardless of its degree of resemblance with the original. Where does this disagreement come from?

To support their argument, Fallis and Mathiesen mention in a footnote the example of a child trying to pass Monopoly money off as real money (call the whole example CANDIES). Such attempt to deceive, they say, cannot also be an attempt to counterfeit money since it lacks a requirement of minimum credibility. Beware that Fallis and Mathiesen do not specify why this case lacks credibility, taking it for granted. Likely, the idea is that the two objects – Monopoly money and real money – do not share a sufficient amount of perceptual similarities that would enable the former to pass off as the latter in an ordinary context[22]. But what if we imagine a different context that makes the trade doable? Here is one.

(4) The sling

Janny wants to persuade with tact his 6 years-old son, Richard, to give her his sling before he hurts himself accidentally. Janny knows that Richard is a very clever child and he has already learned how objects are exchanged in the adult world: by paying. Therefore, the mother hopes to 'buy' the sling from Richard by giving him a pile of Monopoly money passed of a stack of bucks. Janny wants to play on the fact that Richard is too young to distinguish real money from fake money (or even to know the difference).

[22] Notice that this is not the only way through which someone may fail to counterfeit something. Fallis and Mathiesen quote the same excerpt from Gelfert (2018) about the guy from his physics faculty that we have mentioned above, talking about the minimum audience condition (pp. 17-18). In that case, the pseudoscientific paper delivered by the guy was aesthetically equal to a legitimate paper, but the manner of delivery undermined any chance of success in the attempted counterfeiting.

Does Jenny's intent count as counterfeiting? I believe the answer is affirmative, but if the sole criterion for counterfeiting is the one of perceptual similarities the answer should have been negative, for the object that carries the act of counterfeiting is the same as before (Monopoly money). Yet, the judgment for THE SLING is strikingly divergent. Why? To understand this divergence we must first have a preliminary agreement on what it means to counterfeit something. The *Merriam-Webster* dictionary defines the verb "to counterfeit" as to imitate or to feign especially with intent to deceive, or alternatively to make a fraudulent replica of something. We may consider a replica fraudulent when it is capable of pass itself off as the original in the majority of cases and under normal circumstances. Fallis and Mathiesen's use of "counterfeit" seems to be a mix of both alternatives of the *Merriam-Webster* definition: it requires the intention to deceive about its being an original and the propensity to do so; the latter quality being the equivalent of "fraudulent". So in CANDIES, Fallis and Mathiesen would say that the object in question was not fraudulent: this lack does not allow the case to be qualified as an example of counterfeiting, even assuming that there was an intention to deceive on the part of the candy-hungry child.

What about THE SLING instead? If Fallis and Mathiesen want to defend the idea that Janny's is not a real counterfeit attempt, they would likely lean on the peculiarity of the example, arguing that this is not what usually falls under the label of "normal circumstances". Maybe to qualify as a normal context some epistemic requirements on the part of the target must hold. For instance, it is necessary that the target of the counterfeit attempt conceptually knows the difference between an original and a copy of it and is also equipped with the necessary skillset that allows her to actually distinguish them. If so, Richard does not possess the necessary requirements for being the target of a counterfeit attempt. But the problem with THE SLING is that we have a strong intuition that it is indeed a counterfeit attempt. What is missing in the account of Fallis and Mathiesen is then an in-depth explanation of what kinds of context are ordinary ones, given that the notion of propensity (to deceive) always needs to be referred to a model of what we think is an ordinary context or circumstances, and a reasonably justification to dismiss examples like Jenny's as illegitimate.

Quite the opposite for Pritchard (2021), whose position on the matter is peremptory: poorly executed fake news is still fake news. In contrast with Gelfert (2018), who says

that fake news must be «*objectively likely* to mislead its target audience» (p. 108), Pritchard brings the example of a website that is designed to spread inaccurate information – and for this reason it is a fake news source, according to Pritchard's own definition of fake news – whose way of presenting this misinformation as news is so lousy that it is relatively easy for most readers to spot that the news is fake. For Pritchard this is still an example of fake news source which produces fake news as products. Pritchard is explicit on this point, and his position deserves to be quoted in full:

> [T]he fact that the lousy fake news website might not in fact mislead anyone, and certainly wouldn't be such that it is 'objectively likely' that it would mislead its intended audience, doesn't entail that it isn't fake news. [J]ust as an excellent forgery of a Picasso painting does not become the genuine article just because it is so good as to be hard to distinguish from the real thing, so it doesn't follow that a lousy forgery of a Picasso painting is any less of a forgery for the fact that it is easily spotted (Pritchard 2021: 57).[23]

Even though the overall activity of this website is patently ineffective – because it fails to masquerade as news due to its manifest sloppiness – that does not entail that its product is not fake news. But here an opposite problem emerges to that of Fallis and Mathiesen, for Pritchard's position cannot account for our intuitions towards CANDIES. In fact, it seems intuitive to think that the child there failed to counterfeit money, even though there was the intention to do so. The reason, again, is that Monopoly money is not conducive to the success of the counterfeit attempt in that context. However, Pritchard is obliged to commit himself to the opposite conclusion: as much as it was poorly executed, nonetheless the child here has performed an act of counterfeiting.

Consider Pritchard's analogy with a forgery of a Picasso painting. His argument from analogy goes in two directions, both concerning the degree of similarity of the forgery to the original. One is ascending: a forgery that reaches even the highest peaks of imitation does not ipso facto cease to be a forgery. The other one is descending: a forgery so badly made that no one would mistake it for the original does not ipso facto cease to be a forgery. The descending direction of the analogy is what would enable

[23] I take it for granted that by 'mislead its intended audience' he means misleading with regard to the status of the news and not with regard to its content.

Pritchard to say that Monopoly money in CANDIES is a genuine case of counterfeiting[24]. I believe that while the ascending direction is flawless, the descending one is potentially problematic.

To see why, firstly we may run a thought experiment. Imagine a robot equipped with an artificial intelligence programme that enables it to faithfully reproduce all of Van Gogh's paintings. The robot can also reproduce those paintings by changing any of their details if instructed to do so. At the first request, the robot creates from scratch a reproduction of *The Starry Night* which is identical to the original both for material used and for pictorial style. After that, the robot is asked to reproduce the same painting another 200 times, each time changing one detail from the previous one and maintaining this change in the next copy. So for instance, in the second attempt the robot creates a second reproduction of *The Starry Night* that is again identical to the original with the exception of one detail indistinguishable to the naked, untrained eye; say a nuance of colour in a tiny portion of the paint. In the third attempt, the new copy of *The Starry Night* maintains that different nuance of colour and adds a further discrepancy (maybe it changes a bit the position of the moon). The process is repeated for the remaining 198 copies. At an unspecified point along the sequence, the copies of *The Starry Night* no longer resemble the original Van Gogh painting, due to the many discrepancies, and any random observer who does not know robot's instructions could not tell from which painting the robot started. Is every reproduction along the sequence a forged Van Gogh painting?

Pritchard should respond affirmatively, but I doubt this is right. In my view, robot's reproductions are forgeries as long as the robot does not reach that unspecified point where its products do no longer resemble *The Starry Night.* My judgment hinges on the very definition of forgery/counterfeit. A forged x is an object that that is intended to make people believe that it is a genuine X when it is not (Wreen 1980, 2001). As we already said in §7 presenting the account of Fallis and Mathiesen, to be a genuine x is to have the right history of production of x. A forged x is an object that does not have the right history of production of a genuine x but is represented to have it with the intention to deceive (but for now let us leave this intention aside). The first creation of the robot is

[24] 'Forged' and 'counterfeit' are synonymous, since 'forgery' and 'counterfeit' are both logically parasitic concepts on that of genuineness (see Wreen 1980, 2001). Therefore, everything I say about 'forgery' also applies to 'counterfeit' and vice versa.

a forgery, despite its perfect resemblance with the original paintings. The reason is that history of production also entails sameness of the source of issue: the robot's perfect reproduction of a Van Gogh painting is still a forged Van Gogh painting because it did not come out of the hands of Van Gogh. The second creation (the first of the sequence with new instructions) is again a forgery, albeit of a slightly lower quality: it did not come out of the hands of Van Gogh and it is not completely faithful to the original in any aspect. In spite of this further flaw in its history of production, though, it is still a forgery because it is represented to have the right history of production: at first glance and in absence of a further inspection, the average observer would confuse it with real *The Starry Night*. And the same applies to the third and many others until the sequence reaches that unspecified point mentioned above. What happens from now on? Robot's products cease to be forgeries *because they cease to be represented as having the right history of production*. That is, there might be a set of essential properties – like identity with the original subject or pictorial style – that enables a copy of *The Starry Night* to be a genuine copy of that painting and which gets lost during the sequence; even though we cannot say at what point in the sequence they are definitively lost.

And the same may happen with fake news. In the process of becoming less and less similar in appearance to real news, there might be a point of no return after which counterfeit news (i.e. fake news according to privativists) ceases to be counterfeit. If fake news does not have the propensity to deceive due to its similarity with genuine news, maybe it is not fake news after all. Therefore, I find the argument from analogy proffered by Pritchard insufficient to justify the claim that any fake news item, even the furthest from any outward emulation of real news, remains a fake news item[25].

To summarize, my critique is that privativists' thesis of imitation of real news appearance is vulnerable to counterexamples. Privativists either do not specify the properties of contexts in which an attempt to counterfeit is successful – as was the case with Fallis and Mathiesen with THE SLING – or fail to justify the claim that an improbable counterfeit item is still a counterfeit – as for Pritchard with CANDIES.

[25] This is not to say that I disagree with the conclusion. Indeed, as we shall see in the third chapter I actually consider fake news items which do not imitate the appearance of real news as unproblematic instance of fake news. What I do not agree with is the justification Pritchard gives to support this thesis: the idea that a forgery remains always a forgery. Furthermore, I do not think like any other privativists that fake news needs to reproduce the appearance of real news (more on this later).

9. The problem of redundancy, again

In §2 I have summarized the challenges around the definitional project of fake news, starting from the reprimand of Habgood-Coote. The article in which he proposed the elimination of the term "fake news" from the philosophical and journalistic debate dates back to 2019. After its publication, as we know, some scholars criticized his eliminativist solution and endorsed a more tolerant approach. One of the most compelling rebuttals against Habgood-Coote's stance was that fake news differs from other notions employed to talk about miscommunication because it indicates news that imitates the appearance of real news in order to pass itself off as the latter. Imitation of news appearance is a new kind of deception, the argument continues, and that is why fake news deserves to have its place among other philosophical notions.

In a further reply to his detractors, Habgood-Coote (2022) acknowledges that it would be fruitful for the philosophical community to incorporate in its vocabulary a term that denotes the act of deceptive imitation of real news appearance, but he is still sceptical that this could be achieved with the notion of fake news. For one thing, the sense and extension of this imitative action was not fully established by proponents of the privative view. In contrast, there exist an array of different interpretations for the kind of imitation at play. Brown (2019) thinks that fake news mimics real news inasmuch as the former reproduces the conventions of traditional media. But what exactly is imitated by fake news? For Levy (2017), fake news mirrors format and content of legitimate media organizations. The same goes for Rini (2017), whose account of fake news includes the imitation of the conventions of traditional media reportage, but only as a matter of typicality. In contrast, Fallis and Mathiesen (2019) argue that a piece of news must get across the idea that it has gone through the standard modern journalistic processes, while Gelfert (2018) says that fake news mimics the markers of credibility of established news sources[26].

All these nuances point to a lack of uniformity for the imitation thesis: the mirroring of format and content is something different from imitation of standard journalistic practices, which in turn differs from reproduction of the markers of credibility of

[26] Someone may be baffled by the fact that this list of authors includes conservative ones. But here the main question is whether imitation of real news has a univocal interpretation or not. The difference between conservatives and privativists is only that the latter, contrary to the former, believe that this feature is the hallmark of fake news.

traditional media. If we cannot establish what is the unique and legitimate target of the imitative action of fake news, we do not have the right to associate the term fake news with one meaning of imitation at the expense of others. Therefore, Habgood-Coote concludes that we may face a new form of redundancy, for what we come up with is a bundle of new terms, each of which represents a different kind of imitation. For instance, we may have "fake news 1", a term that denotes news which imitates the markers of credibility of traditional media (whatever that means exactly). Then we may coin "fake news 2", which instead is an expression that refers to news which imitates the content and the form of legitimate media organizations; and the list may continue with other examples.

Habgood-Coote's argumentation goes along the following lines. Brown claims that fake news is conceptually necessary inasmuch as it denotes news which imitates the appearance of genuine news. However, there is no legitimate consensus on what is exactly imitated by fake news; there are instead many interpretations of the same act of imitation put forward by different authors. Consequently, there is no decisive reason to associate "fake news" with one specific kind of imitation. At best we can duplicate the "fake news" to account for the multiple nuances of the imitation of real news[27]. At the roots of Habgood-Coote's reasoning is the idea, which I endorse, that without an agreement on what kind of imitation is at play we cannot say that imitation lies at the heart of fake news' conceptual necessity.

Conclusion

This chapter has accomplished three objectives. It presented the challenges around the debate of fake news definition, starting from the provocation of Habgood-Coote of abandoning the term altogether, and it made clear the stance and the purpose of the present work. Then, we have summarized the two main positions that dominate the epistemological debate on fake news: the conservative view and the privative view. Finally, I have gathered the criticism against both positions, and enriched it with my own arguments. My main point is that, beyond counterexamples that target the

[27] I must stress that this argument is only valid as long as the philosophical community does not reach a consensus over this topic. And so far this seems to be the case.

descriptive project of both perspectives, the conservative view and the privative view are essentially unable to solve the problem of conceptual necessity for fake news.

As it has developed so far, the conservative view is perhaps the most exposed to this challenge, since it proposes a concept of fake news that redundantly overlap with other notions (such as mis/disinformation). The privative view gives an innovative meaning to fake news, since the act of deceptive imitation of real news is actually something new. But the problem of redundancy pops up again for lack of uniformity among privativists. As Habgood-Coote states, there is no decisive reason to associate "fake news" with any of these kinds of imitation, because none of them can prevail in legitimacy over the others (in the absence of a strong argument to the contrary).

As discussed in §2.3, the conservative view adopts a narrow understanding of what constitutes news – and by extension fake news. Within this framework, news is merely a type of information obtained from a news source, effectively equating the concept of news with that of information (differing, at most, in terms of their dictionary definitions). The purpose of the next chapter is to demonstrate why this equivalence is flawed and why it must be reconsidered before proposing a new definition of fake news.

Chapter 2

The nature of news

1. Searching for the culprit

In the first chapter I presented the debate on the definition of fake news as seized by two competing models, the conservative view and the privative view. The former assumes that fake news is a kind of news, while the latter holds the opposite: fake news is not news. I pointed out that both views propose a redundant conception of fake news, although in different ways. Conservatives often endorse a notion of fake news that ends up overlapping with already existing notions that capture pathologies of information – such as misinformation and disinformation, much discussed in philosophy and communication science. On the other hand, while not suffering from the same vulnerability, the privative view fails to establish once and for all which aspect of real news is imitated by fake news. Therefore, the redundancy stems from the fact that we should coin a new term "fake news" for every privative account which associates it with a different sense of imitation[28]. In either case, the current positions on fake news incur a problem of conceptual necessity. Is there a way out from this?

This chapter considers whether the root of the problem lies in how conservative and privative authors define the very notion of news. Indeed, I argue that the main issue of prevailing accounts is that their notion of news is excessively minimalist, and most of the time it ends up coinciding with some kind of *information* – and specifically with a non-factive interpretation of information (viz. information that can be either true or false). This assumption is troubling, since treating news as a kind of information directly influences the way in which conservatives and privativists deal with fake news. Instead, I maintain that news is qualitatively different from information, and identifying the former with the latter is a blunder. While information is a semantic entity, news in my view is a sociological status that a piece of information inherits from a process of verification and elaboration by specific institutions; and these are journalistic sources.

[28] This is not the only reason why, in my advice, the privative view is wrong about imitation. Fake news does not necessarily imitate the appearance of real news. Indeed, we can deem something as fake news even though it is outwardly divergent from what we normally think of as the look of real news. The fundamental reason is that there is no official set of outward characteristics of a news item. There is no decalogue, no universally recognised list of superficial elements that a news item must necessarily possess in order to be 'real' news. What we have instead is a set of flexible expectations that each time enables us to accommodate our judgment on news items in front of us. More on this below, in §6.2.

In its several expressions that range from printed media to digital ones, journalism is the social institution that has the monopoly on the production and the publication of news. Thus, a news item is not just a semantic entity. It is instead an institutional entity inasmuch as its existence depends to the observation of some settled journalistic norms that characterize its process of production and publication. What we needs in order to understand the essence of news, is a general theory that explains what is an institution and how sociological entities arise from institutions on which they depend. I believe that Searle (1995, 2005, 2010)'s theory on generation, maintenance and destruction of institutional facts is adequate for the purpose.

The objective of this chapter is twofold. First of all, I intend to show where the problem of conceptual necessity for the conservative view of fake news is located (§2). As partially anticipated, this is all due to an overly minimalist interpretation of what news is (§3). Secondly, I explain in what sense news is a status, starting from the interpretation given by the privative view (§4). However, we will see that this interpretation is only partially correct and can be reformulated by drawing on Searle's theory of the construction of social reality (§5). Once the conceptual tools of Searle's theory are adopted, I proceed to give my personal interpretation of what news is (§6), and then explain in more detail what is wrong with the proposal of the privative view.

2. Prevailing positions on news

If we were to look for a definition of news within the broader epistemological debate on fake news, we wouldn't probably find anything conclusive on the subject. Sometimes scholars do not even make such a big deal out of it, and go on with their analysis of fake news without spending too many words on what news is. Other times a definition of news can be derived only indirectly from the philosophical account at issue. For example, it is quite clear that for Mukerji (2018) a news item is propositional content asserted or implied in the form of news publication, even though the author never states it explicitly. Mukerji's putative definition may strike as a bit circular, since it makes reference to the notion of news («form of news publication») within the definition of news itself. This is not the real drawback of his proposal though. As he keenly points out, talking about the form of news publication exposes his account to a problem of vagueness. How can we trace a clear boundary between what has the form of news

publication and what does not? We may make a list of possible distribution channels for news publications and say, for instance, that anything published on newspapers or news websites has the form of news publication. But there is a whole range of borderline cases that are difficult to classify – as Mukerji himself acknowledges. Consider for instance the example of a BBC journalist who posts a story to a private social media account that has a substantial following. Is it a story in the form of news publication or not? Different people will have different intuitions on this case, and there is no univocal assessment.

Authors belonging to the conservative view – who cover the majority of the proposals in the field – usually endorse a minimalist account of news as propositional content of some kind. Most of the time it is asserted content, but not necessarily. It must be emphasised, though, that assertion here is not always the same as we find in the literature on speech acts (for an overview, see Pagin and Marsili 2021). It is instead a looser concept of assertion for at least two reasons. For one thing, here not only utterances can be assertions but also any communicative act by which a propositional content can be conveyed: photos, videos and any message transmitted in multimedia form can count as assertions under certain conditions. Secondly, in the literature on speech acts, as a rule, asserted content is always distinct from merely implicated content. This is not always the case for conservative authors here, who sometimes consider merely implicated content as part and parcel of the asserted content.

For the privative view the question is much simpler, because here the definition of fake news arise by contrast from the one of genuine news. For many privativists, then, genuine news is information or stories that have gone through standard journalistic processes (one exception is Pritchard, p. 44 above, for whom news is just information that someone gets from a news source, and therefore the question for him becomes "what is a news source?"). According to the privative view, then, news is not a thing but a status: a property that supervenes on a piece of information if certain conditions are met, where that conditions are the fulfilment of standard journalistic processes.

My plan for the next sections is to critically discuss the idea that news is reducible to propositional content of some kind. I believe that this conception either does not match our intuitions around the issue or runs again into the problem of redundancy (§3.1). Instead, I maintain that the general idea of the privative view is true: news is a status

that information gets under certain conditions (§4). However, this status is not straightforwardly acquired the way privativists say. I will elaborate more on this in §6 and beyond.

3. News as reported information

Almost every conservative account holds that news is propositional content of some kind. The view of news as propositional content stems from a minimalist notion of news as information. For instance, in developing their own definition of news Jaster and Lanius start from the *Oxford English Dictionary* entry:

> In our account, we are employing the *Oxford English Dictionary*'s (2018) minimalist notion of "news" as "newly received or noteworthy information, especially about recent events" and, in a more specific usage, as a "broadcast or published report of news". Accordingly, we use "news" to refer to any report of typically recent events that is broadcast by media or individuals to address a public. (Jaster and Lanius 2021: 20)

News here is noteworthy information about recent events plus a bunch of other properties, such as a feature of attributability (the piece of information can be credited to the activity of a news source, whether it is an institution or an individual) and a naïve version of the minimum audience condition («to address a public»)[29]. But what is information from a philosophical perspective?

We cannot dig too much in the complex philosophical literature on information here. Suffice it to say that whenever we talk about information transmitted through either linguistic or extralinguistic devices we are dealing with well-formed and meaningful data. This characterization revolves around the notion of data, which is a very elusive one. Floridi (2010) suggests that a good way to grasp the nature of data is understanding what it means to lose them. Here is his example. Imagine a page from a book written in an unfamiliar language. The data are presented as pictograms, arranged in regular patterns that suggest adherence to some structural rules. Although we possess all the data, their meaning remains unknown, and thus, we lack any information. Now,

[29] As we shall see, the stance of Jaster and Lanius on the matter is more complex since they treat news as asserted content. Nevertheless, their starting point is the idea that news is information, for to assert something there must be something, i.e. propositional content, that can be asserted to begin with: «in the contest of fake news, "information" [...] needs to refer to the communicative *content* of utterances» (Jaster and Lanius 2021: 36).

consider if we were to erase half of these pictograms. It could be argued that we have also reduced the amount of data by half. If we continue this process until only one pictogram remains, we might be inclined to suggest that data coincide with certain representations. However, if we were to erase that last pictogram as well, we would be left with a blank page. Nevertheless, even a blank page is not entirely devoid of data. The presence of a blank page remains a datum, as long as there exists a distinction between the blank page and one that contains or could potentially contain written content. Floridi notes that here there is an analogy with silence. Even someone who keeps quiet can potentially release data, if lack of perceivable data stands for other data – as when two interlocutors agree to assign a specific meaning to silence with the goal of being undetected. The reality is that a true, total erasure of all data can only be accomplished by eliminating all conceivable distinctions. This explains why a datum is fundamentally reducible to an absence of uniformity. Consequently, Floridi offers the following definition of datum:

> DATUM = *x* being distinct from *y*, where x and y are two uninterpreted variables and the relation of 'being distinct', as well as the domain, are left open to further interpretation.

Lack of uniformity can be found among physical states/signals as well as among two symbols of a written language.

It is not enough to have data to have information, though. Data must possess two further properties. In order to be information data must be well-formed and meaningful. Data are well-formed when they are arranged according to the rules of the code or the language through which they are displaced. In a nutshell, to be well-formed is to observe the syntax of a given language. For example "$) \neg (x \rightarrow \forall a)$" fails to be a well-formed formula of first-order logic since it does not respect the syntactic rules of that language. Finally, data must be also meaningful: they need to comply with the semantics of the language in question. For instance, in the language of classic first-order logic the symbol "□" has no meaning, whereas it does in modal logic (it is a modal operator that stands for "it is necessary that…")[30].

[30] It must be emphasized that data constituting information can be meaningful even in the absence of a competent interpreter of the language in question. The absence of an interpreter is a matter of information accessibility, and it does not impinge upon the semantics of information.

Therefore, information is well-formed data that possesses a semantic content, and thus represents some part of the world as being a certain way. Information is a representational entity (Fallis 2015). At this point, theorists of information usually split into two camps. There are those who embrace a factive account of information (Dretske 1981, 1983; Floridi 2005, 2010, 2011; Dinneen and Brauner 2015) according to whom information must be necessarily veridical semantic content[31]. That information is factive means that information is necessarily veridical. Due to this feature, theorists who endorse a factive interpretation of information think that information has a structural relation with knowledge. Knowledge encapsulates veridicality because it encapsulates semantic information, which in turn encapsulates veridicality. False information (aka misinformation), for these theorists, is just a contradiction: false information is not information at all. Other theorists, instead, believe that information is non-factive: it can be as veridical as false since it is a purely semantic entity that bears no essential relation with truth (Scarantino & Piccinini 2010; Cevolani 2011; D'Alfonso 2011). As Simion (2023) notes, for the second group information here is almost the same as meaning: the headline "Pope Francis endorses Donald Trump" conveys the false information that pope Francis endorses Donald Trump in virtue of the fact that it means that pope Francis endorses Donald Trump.

Regarding news, conservative authors have in part a similar position to that of non-factive theorists of information. As a first step, news is like non-factive information in the sense that it is propositional content conveyed in many forms (such as articles, photos, videos) which can be either true or false. News does not cease to be such if it lacks (only) truth – just as false information remains information for non-factive accounts. Then, many conservatives add a further step[32]. According to them, when we talk about news we are not considering just the propositional content conveyed by a news item, but also the very act of reporting it. That is, news is propositional content asserted by someone with specific intentions and in a specific setting. The intention is to get the addressees accurately informed about relevant facts, while the setting is the one

[31] We use "veridical" in the place of "true" because strings or patterns of well-formed and meaningful data can constitute sentences in a natural language; however, they can also generate formulas, maps, diagrams, videos, or other semiotic constructs in a variety of physical codes. In these cases, "veridical" is preferred over "true" (Floridi 2010: 50).

[32] Not all conservative authors adhere to the idea that news is asserted content. Rini (2017), for instance, stops at the first step in saying that news is a kind information since fake news is primarily something that «purports to describe events in the real world» (p. E-45), albeit poorly.

of news media, in which someone (journalists, editors etc.) is addressing an audience composed of more than one person. Therefore, news is not just information but information reported (i.e. asserted) under particular conditions and circumstances.

Many conservatives clearly agree on this point. Jaster and Lanius specify that their definition of news can be tweaked into a more normative one, since they conceive reports as truth-warranting assertions (Carson 2006, 2010). According to this view, assertions are something along the lines of promises: with an assertion, the speaker guarantees either explicitly or implicitly that what she says is true. They are speech acts by which the speaker warrants the truth of the propositional content of her utterance (yet it should not be forgotten that it is only one of the many possible theories on assertions offered by the literature. Jaster and Lanius simply considered this to be the interpretation of assertions most compatible/adequate with their account of fake news, and the same goes below).

Even though his position is not fully clear, also Gelfert (2018) seems to be converging towards a concept of news as asserted content, whilst he defines news as claims *presented to us as true* by a putative news source (p. 104). The idea that to assert a content means presenting it as true has a long tradition and comes from authors like Frege (189271948), Wright (1992), and Adler (2002). Marsili (2024) presents the core of this account on assertions by means of the following comparison. An unasserted proposition simply describes or represents a state of affairs without taking a stance on whether that state of affairs corresponds to reality. Here we are still at the first step, that of pure propositional content (or non-factive information, in our lexicon). Such a proposition can be deemed true or false, but not correct or incorrect. This is because a proposition, on its own, does not provide a criterion to evaluate its correctness; it merely represents a state of affairs without specifying whether that state actually holds. On the other hand, asserting a proposition involves presenting it as true – describing the world as being a certain way – so that the assertion is successful (correct) only if the world is indeed that way. News is then information presented to a public as true by news sources. The same position is shared by Croce and Piazza (2022), although they, like Gelfert, do not commit themselves strongly to any theoretical account of assertions.

Despite being so popular among conservative accounts, in my view, the collapse of the notion of news over that of information is at the heart of the problem of conceptual

necessity for fake news. I shall also argue that the conception of news as merely reported information clashes with our intuitions around ordinary cases of news. I will extensively develop this point in the next two subsections.

3.1 News is not just information

In the first chapter I said that conservative accounts of fake news run into a problem of conceptual necessity. This issue points to the fact that there are other, more clear notions – such as misinformation, disinformation, bullshit – which already do the work that "fake news" is supposed to do. Here I have been claiming that the source of this problem may lie in the way the notion of news is cashed out.

For conservative authors, it seems that news is reported information, whereas "reported" indicates both the performance of an assertion and the fulfilment of some extrinsic conditions (recentness of information and minimum audience condition), and information is interpreted as non-factive. Here is the scheme for the conservative (c) conception of news:

> $NEWS_c$ = recent asserted information transmitted to a minimum number of consumers

Under this paradigm, properties of fake news such as lack of truth and/or lack of truthfulness are added to this basis. Fake news is then a kind of news with additional epistemically-bad features. Here is a possible scheme for fake news in the account of Jaster and Lanius (2021) (jl):

> FAKE $NEWS_{jl}$ = recent asserted false and/or misleading information transmitted to a minimum number of consumers with the intention to deceive them or with an attitude of indifference towards the truth

Even non-hybrid accounts follow the same pattern and can be reconstructed starting from the same basis. Below the possible scheme for the account of Mukerji (2018) (m):

FAKE $NEWS_m$ = recent asserted information transmitted with a covered-up attitude of indifference towards the truth[33]

And, finally, the same goes for the minoritarian positions that do not include the assertion part in ($NEWS_C$). One of them is the account of Rini (r):

FAKE $NEWS_r$ = false information transmitted to a minimum number of consumers with the intention to deceive them

What stays constant in any conservative proposal is the fact that news is reducible to some sort of information. But here we face the challenge of conceptual necessity presented in the first chapter.

Since there are many other concepts that already describe pathologies of information, fake news becomes redundant because any conservative definition is just the repetition of those concepts with the addition of ad hoc properties. Someone may disagree by pointing to the fact that news in many conservatives account is not just information, but reported information; that is, we should see it more under the lens of the concept of assertion than that of information. Although this is reasonable, notice that this shift does not solve at all the problem of redundancy. Indeed, it may exacerbate the issue.

Recall that the challenge of conceptual necessity consists in the fact that we can analyse the predicament we find ourselves in right now – e.g. the quality of information on social media sites – without resorting to any new term. The reason is that we have plenty of concepts for talking about deceit, miscommunication and epistemic dysfunction (Habgood-Coote 2019: 15). And it does seem that assertion-based definitions of fake news are actually reducible to those pre-existing concepts.

Take the example of FAKE $NEWS_{jl}$. Jaster and Lanius talk about news as asserted content, and then use the properties of "lacking truth" and "lacking truthfulness" to group together different cases. A news item lacks truth when it is either false or misleading, while it is devoid of truthfulness when it is published with the intent to deceive or without caring about the truth. If we plot all possible combinations we obtain the following four cases of fake news:

[33] Mukerji does not defend any minimum audience condition for news, but he only talks about assertions made with the form of news publication.

1) False assertion with the intent to deceive

2) False assertion with indifference towards the truth

3) Misleading assertion with the intent to deceive

4) Misleading assertion with indifference towards the truth

Note that every case it is simply the juxtaposition of three well-known notions in different combinations: false/misleading assertion, bullshit, intention to deceive. Alternatively, one can even use slightly more complex notions that are suitable for some cases. For instance, we said that Jaster and Lanius endorse the idea from Carson (2006, 2010) according to which assertions are truth-warranting utterances. Carson himself uses this interpretation to construe his own definition of lying as a kind of assertion:

> A person S tells a lie to a person S1 iff:
> 1. S makes a false statement x to S1;
> 2. S believes that x is false or probably false or S does not believe that x is true;
> 3. S states x in a context in which S thereby warrants the truth of x to S1;
> 4. S does not take herself to be not warranting the truth of what she says to S1
> (Carson 2006: 298)

Notice that the first point of Carson's definition requires the statement of a lie to be false. Therefore the first case from the above list of combinations can just be treated as a lie with an intent to deceive, according to the same conception of assertions accepted by Jaster and Lanius. Nothing exceptionally new. In any case, we do not need to resort to any new concept in order to account for cases (1)-(4). Existing notions in philosophy of language and pragmatics are largely sufficient. What would be the point of a neologism?

Again, it could be argued that even though "fake news" can be decomposed into other notions, including assertion, there is however an aspect of novelty: they would in fact be assertions in the form of news. This reply follows the same argument about bullshit already addressed in the previous chapter (§6). The problem with it is that, as long as being in the form of news only means either addressing a public or being transmitted by a news source, the lion's share in fake news definition is performed uniquely by the concept of assertion. Therefore, we do not have sufficient reasons to say

that the addition of “in the form of news” justifies the coining of a new notion from scratch. We are then back to square one: how can we fix this?

In my view, the key move is to give up the idea that news is reducible to non-factive information reported by someone. In fact, I believe that this is an unduly minimalist conception of news which unfortunately gets validated by idiomatic uses of the term “news” in ordinary conversations. For instance, sometimes we use sentences like “she has some news” or “you gave me good news” without thinking twice, thus legitimating the equivalence between news and reported information. But it is also true that here we are disputing the notion of news within a highly-specialized philosophical framework: we cannot settle for the naive interpretation conveyed by an idiomatic use. To see why we may discard the idea of news as reported information, we may test its validity through examples until we reach a point of no return.

Suppose that friend of yours texts you saying that she saw a terrible car accident along the street that passes in front of your house. Later that day, your local newscast transmits a detailed report on the same event that reaches an enormous amount of people and arouses clamour in your neighbourhood. Are we entitled to say that the piece of information reported by your acquaintance is the same thing as the newscast report? Most likely, our intuitions would suggest a negative answer. However, a conservative author may reply that the issue here is just that the piece of information reported by your friend is told in private; unlike the newscast report, it does not count as news because it does not address the public. So the negative answer, they continue, is not driven by a rejection of the equivalence between news and reported information. It is just that the report of your friend could not reach a sufficient amount of people.

Let’s accept for the time being this reply for the sake of the argument, and go on by tweaking a little bit the example. Consider the same scenario except your friend communicates the fact via social media. She writes a public post on Facebook recounting the incident and offering more or less the same amount of information that she shared via text in the first scenario. The post is public and therefore it is available to a large audience. Now we ask ourselves the same question: is the piece of information reported by your acquaintance on social media the same thing as the newscast report? Again, a newscast report seems at first to be qualitatively different from the Facebook post of your friend, even though I may not be able to spell out why at first. This time the

conservative author cannot appeal to some version of the minimum audience condition to justify my intuitions, since the Facebook post is there for everyone to read (provided that privacy restrictions are not in force). It could be argued that the difference lies in the format: the Facebook post does not have the news format as the newscast does. But what is the news format? Many conservative authors either adopt an ostensive definition – the news format is the format of articles, photos and videos published by news sources – or make it coincide with some version of the minimum audience condition. We have already discarded the latter with the last example, so we may test the ostensive definition. Let's tweak again our example!

Imagine again the same initial scenario, but this time your friend buys a WordPress domain to create a temporary information website where she can publish and disseminate news relevant to her and her local community. Your friend is not very competent and is also hasty in gathering information, but she thinks that in the end the most important thing is to report the fact. On this occasion, then, she writes an article that reports what she saw of the car accident, including the testimony of some passers-by and police officers who reached the location. She then uploads the article on the Facebook page of her website and so her report is both public and in the form of news article. Again, is this the same as the evening newscast report on the car accident? I doubt it, since it is not enough to imitate the appearance of a thing to be that thing. In the same way, I have the strong intuition that having the form of news does not necessarily make a piece of information news; an intuition that is shared also by privative authors, by the way.

At this point, even the conservative author may agree on the fact that the main difference between the two cases is that the newscast, unlike the text/post/personal website, is the product of the activity of a specific institution. The essential difference that I want to stress is that information encapsulated in the newscast report seems to have an additional value. Unlike the testimony or the post of your friend, the newscast report is in part the outcome of an activity of selection and verification that makes it more reliable and valuable overall. And this is because it was made by an institution (or by one of its representatives) *we trust*. It therefore seems that when we use the term "news" we are referring to an object that has the property of being produced and

distributed by a reliable institution, and not to a semantic content of some kind. Which brings me to the second conception of news available in the market: news as a status.

4. News as a status

One of the contributions of the privative view is to set the problem of fake news as a problem of status. The idea is simple. Even though the expression itself may fool us, we should not think that fake news is a kind of news. Instead, fake news does not really possess the status of news, but it only gives us the impression of having it. Fake news is something that masquerades as news for some purpose, like deceiving someone. The great takeaway here is that news is not set out as an object of some kind. To be news does not mean to be some sort of asserted content, nor any other thing. To be news is instead to possess a specific status: the one of "news". Then the question becomes "what is the status of news?".

Privativists usually believe that the possession of this status hinges on *having the right history of production of news*. That is, for something being news – be it a piece of information or a story[34] – it must have gone through the same production process that every news item normally undergoes. The description of this process may vary across privative accounts, but its essence remains the same. Basically, it is the process of verification and double-checking executed by those who deal with news by profession. It is the process of news creation carried out by journalists. In sum, something that has the appearance of news is really news only if it has been produced according to the standard of journalistic production.

Privativists use the expression "genuine news" to designate objects that have the appearance of news *and* have been produced according to journalistic standards:

> By genuine news, we mean stories that have gone through the standard modern journalistic process […]. That is, genuine news has been produced by professionally trained reporters, fact checkers, and editors, who are attempting to provide fair and accurate accounts of current events. (Fallis and Mathiesen 2019: 8)

[34] I rather avoid to use the term "story", since I find it unsuitable. It is a term that is too related to the idea of narration, even though a news item is not always the narration of something. The daily news bulletin about the stock market does not narrate anything *stricto sensu*, and yet it is news for all intents and purposes. Henceforth, I will limit myself to using only the term "information", even though many privativists freely use "story".

“Genuine” is usually regarded as subsective adjective (Cappelle et al. 2018, Partee 2021). Subsective adjectives, like “skillful” or “good”, are so called because the denotation of a subsective adjective-noun combination is a subset of the denotation of the noun. Contrary to other adjectives, like “red” or “carnivorous”, subsective have a tighter semantic relation with the meaning of the noun they apply to. If we know that Peter is both a tennis player and a pianist, from the assumption that Peter is a skillful tennis player we cannot automatically conclude that Peter is also a skillful pianist. On the other hand, If I say (truthfully) that Tasmanian devils are carnivorous animals, I can equally say that Tasmanian devils are carnivorous marsupials without thereby changing the truth-value of the sentence.

Within this model, “genuine” behaves more like “skillful”. If we know that an old document is both an historical find and a birth certificate, from the assumption that the same document is a genuine historical find we cannot say for sure that it is also a genuine birth certificate of the time; it may well be a forged birth certificate, as much as being an historical discovery tells us interesting details about the making of forgeries in the past. Below is the characteristic inference pattern of subsectives (where “#” denotes a wrong inference):

1a) X is a genuine historical find > X is an historical find

1b) X is a genuine historical find > # X is a genuine birth certificate

Notice that while subsective adjectives cannot be automatically transferred to new nominal phrases salva veritate (1b), they nevertheless allow the following inference: if X is Adj(Np) then X is Np (1a). If a document is a genuine historical find, then the document is an historical find. In the same way, if a piece of information is genuine news, then the same piece of information is news.

It is the other way around in the case of fake news. “Fake news”, for privativists, designates something that has the appearance of news but at the same time does not have the right history of production. And for this very reason, it does not possess the status of news. Within the privative model, the adjective “fake” in “fake news” is interpreted as a full-fledged privative adjective. Like “artificial”, “false”, “fictional” and many others, privative adjectives are words that contrary to other categories of adjectives do not describe or say something more about an object. Instead, they usually

deny that the referent of the noun phrase (in this case, "fake news") belongs to the extension of the noun to which they are attached ("news")[35]. In other words, privative adjectives do not permit inferences like (1a), and for this very reason they are classified as non-subsective adjectives together with modal adjectives ("alleged", "possible" etc.). Defenders of the privative view interpret the expression "fake news" along the lines of "fake gun". Just like a fake gun is an object that looks like a gun but it is not actually a gun, a fake news item is something that looks like a genuine news item when actually is not. In both cases, the inference that goes from "fake x" to "x" is not allowed. The privative view may resort to this indirect and defeasible evidence from experimental linguistics to strengthen its interpretation of fake news. If the usual behaviour of the adjective "fake" is that of a privative adjective as in (2), then we may be entitled to think that the same happens with news: fake news is not news.

But why does having the right history of production matter? What does it add to a piece of information? After all, we know that the observance of journalistic standards cannot always ensure the veridicality of a news item. So why bother so much with it? The answer of the privative view is that having the right history of production bestows an *intrinsic value* upon the piece of information in question.

Lazer et al. (2018) suggested that news media's editorial norms and processes ensure the accuracy and credibility of information. A piece of information that has been checked and verified according to journalistic standards has therefore the mark of accuracy and credibility. Other privative authors go as far as to say that the intrinsic value in question is *epistemic*. Pepp, Michaelson and Sterken are explicit on the point:

> [W]hen a story is treated as the product of standard journalistic practices, it is treated as giving a certain kind of *epistemic status* to the information it contains [emphasis mine]. This is not merely a certain level of reliability, but also that the reliability itself is transparent to a significant degree. Journalism lets us understand not only how things are but why we should believe that they are that way. (Pepp et al. 2019: 78; footnote 21) [36]

[35] I say "usually" because, as we shall see, this is not always the case with some nouns. Sometimes privative adjective perform a different job.

[36] The excerpt deserves a note on terminology. Pepp, Michaelson and Sterken talk about stories «treated as» the product of standard journalistic practices because of their own definition of fake news, which assumes the point of view of news propagators (see p. 37 above for the distinction among producers, consumers, receivers and propagators of fake news).

Their article goes on elucidating the epistemology of journalism, which consists in the reliance on information obtained non-inferentially, through either testimony or photographic/video evidence, and the transparent disclosure of the sources (at least as long as the source does not need to be protected).

Despite his aversion towards grounding the status of news on having the right history of production, also Pritchard believes that the observance of journalistic standards is a sign of good epistemic practices (above, p. 44)[37]. Checking sources, corroborating information, seeking confirmation or disconfirmation of one's information… these practices, as many others, «have an epistemic import» (Pritchard 2021: 54). Thus, news sources which follow journalistic standards are epistemically good and are more reliable, ceteris paribus, than news sources that adopt poor practices.

Finally, Fallis and Mathiesen (2019) say that the epistemic value is the reason why history of production matters more for genuine news than for other objects that can be counterfeit, like coins. Take the example of United States currency. That some bits of paper and metal produced by the United States Mint count as currency with legal tender in the context of United States is a fact made possible by collective agreement, according to Fallis and Mathiesen. Buyers and sellers in the United States collectively treat those bits as valuable and for this very reason they are valuable. Their value does not depend too much on their history of production, they continue. In contrast, a genuine news item has epistemic value whether or not readers treat it as having one, precisely because epistemic value stems from the work of professional journalists that produced that news item in the first place; and this work consists in gathering evidence and checking for mistakes. Therefore the history of production of news items matters a lot.

On the basis of this observation, Fallis and Mathiesen suggest that, in order to describe news value, a better analogy would be between fake news and counterfeit versions of brands that are valued for their quality. Fake news are like fake Gucci handbags. Gucci handbags are produced through a process that pays attention to the quality of raw materials (like leather), the efficiency of the treatments of the same materials and the professionalism of the artisans who work on the product throughout the whole supply chain. The output of such a process is then a commodity of a superior

[37] Pritchard explicitly says that «news is a kind of information that someone gets from a news source» (2021: 52). Therefore, his account is among those criticized in the previous paragraph.

standard. Fake Gucci handbags mimic the outward features of genuine Gucci bags to pass themselves off as real Gucci products of superior quality. In the same manner, fake news mimics the appearance of news to pass themselves off as object with epistemic value. In both cases, though, the intrinsic value of the genuine tokens exists regardless of how people collectively behave. The enhanced quality of a Gucci handbag endures even though buyers abruptly think differently. And so it goes for news and its readers.

To sum it up, the privative view supports a conception of news as a form of status. According to this idea, news is not a kind of information but a sort of quality mark that we attach onto it; in the same way as we label Gucci or Dior as quality brands. To get this mark, a piece of information must possess the property of having the right history of production. Specifically, it must have gone through the process of production and verification that meets the standards set out by journalists. The observance of the journalistic standards should guarantee the intrinsic epistemic value of the piece of information, since those norms are geared towards maximising the accuracy and credibility of that same piece of information.

Generally speaking, I agree with the privative view that news is more like a property that supervenes on information, rather than just a species of information. I believe that treating news as a status is correct, although I harbour reservations about some important aspects of the privative account. In order to present these objections, however, it is first necessary to make a digression on what a status is and how it can be attributed, since the privative view is defective on this point. My thesis is that the collective behaviours of news consumers does in fact play a role in bestowing the status of news, contrary to what has been said before. Furthermore, and this is the main point of disagreement with the privative view, I do believe that in explaining what fake news is we do not need to resort to any form of imitation of the appearance of real news. Indeed, a news item can safely be fake news without reproducing the aspect of genuine news, the reason being that *there is no accepted universal canon for the appearance of news*. Instead, what I will try to argue is that anyone nurtures a set of expectations about news appearance which are, however, balanced by a subjective element of evaluation. What is genuine news for me, then, may not be always genuine news for someone else. But this is not to say that there is absolute relativism on the matter: it is still possible to formulate a reasonable list of rules that should dictate the way in which news items are

designed and presented to the public audience. Nonetheless, it is not a given neither that news consumers will recognise those rules as universally valid nor that journalists will interpret and apply them in the same way.

In the rest of this chapter, I will set the foundation of my account on fake news by presenting my own conception of news as a special class of institutional entities. I barrow this expression from John Searle's theory of institutional facts, which are facts that exist only in virtue of some antecedent human institution. Here I speak about entities rather than facts because I find it more perspicuous when it comes to news, although the principle is the same: a news item is an entity that depends on the prior existence of a human institution. To understand this, we then need to explain what this dependency mechanism is based on and how generation of news works. Therefore, in the next paragraphs I will summarise the core principle of Searle's theory that may serve us in the subsequent analysis of news. After this long detour, I will proceed to examine what I think is still wrong with the privative view, and this has mainly to do with the role of collective behaviours in determining the value of news and the idea of imitation.

5. What is an institution?

In his groundbreaking work *The Construction of Social Reality* (1995), Searle provides a theory of how peculiar social facts like private property, money and marriages can arise from sheer physical reality. The philosopher deems it as a critical question inasmuch as, according to him, only human beings are capable of creating such a continuity from the physical world, composed of atoms and particles, to our complex social world, constituted by institutions that are not found in other animal species. For the purpose of our work, it is sufficient to recap only those arguments from Searle's work which enable us to reconstruct his conception of institutional facts. Searle's strategy is to answer the question "What is an institution?" by explaining what is an institutional fact first. Then, he generalizes on institutional facts to get an idea of what institutions are. In my summary of Searle's theory, I will follow his own argumentative structure: first institutional facts (or entities), then institutions.

Searle starts his investigation on institutions by making some theoretical distinctions, the first being the one between observer independent features of the world and observer

relative ones. The former include things like force, mass, gravitational attraction, chemical bonds, and tectonic plates and so on; the latter include other things like money, government, presidential elections, and football games. In a nutshell, observer independent features (OIF) are such because they exist independently of our representations of them; they are intrinsic to nature, and are given also in the absence of an external observer. It is the other way around in the case of observer relative features (ORF), which exist only relative to the attitudes and the intentionality of an observer. ORFs do not add new material objects to reality; they add instead new descriptions of the same object. For instance, I can say at the same time, and without contradiction, that the wooden object in front of me has a certain weight and mass (which are OIFs) and that it is a rolling pin (an ORF). Being a rolling pin is having a feature that depends on the fact that people agree to use that object as a rolling pin. Searle says that a rapid test to understand whether or not a feature is OIF or ORF is: could the feature have existed if there had never been any conscious human beings with any intentional states?

A second distinction, which intersects with the first, is between an epistemic sense of objectivity and subjectivity and an ontological sense of this pair. Epistemic objectivity and subjectivity are features of statements. For instance, If I say "The Starry Night was made by Van Gogh" I claim something epistemically objective because my statement can be established as true or false independently of the attitudes and the feelings of the utterer (me) or other participants in the discourse. But if I say "The Starry Night is Van Gogh's most beautiful painting" I state something epistemically subjective, since its truth is a matter of opinion and personal tastes. The second sense of the objective-subjective distinction is ontological. Some phenomena exist only as intimate experiences of human and non-human animals: pains, itches, mental visualizations and so on. They are ontologically subjective entities because they cannot exist outside the experience of a subject (human or otherwise). Mountains, trees and stars do not need instead the intimate experience of a subject in order to exist; they exist regardless of the inner reality of an animal being.

These two distinctions serve the purpose of framing the problem of Searle's dissertation: how can there be an epistemically objective institutional reality, which includes money, property and other stuff, given that this reality is in part constituted by

subjective feelings and attitudes, viz. ontologically subjective entities?[38] In other words, how can we accommodate the fact that we can say truthfully that two people are married despite this very fact is made possible by a combination of feelings and intentions? To make a long story short, Searle believes that we may solve this puzzle by resorting to three notions: collective intentionality, the assignment of function by human agreement, and the concept of status function. Let's see them in brief as presented by Searle.

5.1 Collective intentionality

The first concept we need to consider is collective intentionality. "Intentionality" is a term used by Searle to describe the feature of the mind that is directed towards objects and states of affairs in the world, such as beliefs, hopes, fears, desires, and emotions. In this technical sense, he argues, intentionality does not necessarily imply any connection with intending in the everyday sense, like intending to go to the movies tonight. Instead, intentionality is a broad concept concerning the directedness of the mind. Ordinary intending is just a specific instance of intentionality, alongside beliefs, desires, hopes, fears, love, hate, pride, shame, perception, disgust, and others.

After having established this, Searle notices that a notable characteristic of humans and many other animal species is that they can engage in cooperative behaviours. Examples include playing in an orchestra, participating in team sports, or simply having a conversation. In such cases, individuals act independently, but their individual actions – playing a violin part, passing a ball to another player – are performed as part of collective behaviour. During collective action, an agent performs their actions as part of a larger group effort. Therefore, collective intentionality is the power of minds to be *jointly directed* at objects, matters of fact, states of affairs, goals, or values (Schweikard and Schmid 2020).

Searle believes that collective intentionality forms the foundation of all societies, whether human or animal. Humans, like many animal species, possess the capacity for collective intentionality and thereby the ability to form societies. According to Searle, though, we should not make the mistake to reduce collective intentionality to individual

[38] This is the original question that drives the inquiry in the 1995 book. However, Searle (2005) subsequently acknowledges that the question is wrong since there are institutional entities, such as corporations, that exist totally independently of physical properties (Smith & Searle 2003).

intentionalities plus a constrain on mutual beliefs. The idea would be that if we intend to do something together, then that consists in the fact that I intend to do it and I also believe that you intend to do it too; and you intend to do it and you also believe that I intend to do it too; and so on ad infinitum for all other agents involved. This reductionism is mistaken and is proven to have many vulnerabilities (Searle 1990). So Searle discards this account and embraces the idea that collective intentionality is a biologically primitive phenomenon.

Afterwards, by stipulation, Searle defines a *social fact* as any fact involving the collective intentionality of two or more agents and *institutional facts* as a subclass of social facts which involves human institutions. Two people that are playing football together is then a social fact; whereas the fact that the same two people are married is an institutional fact, because it involves the prior existence of the institution of marriage in order to exist as well. These are the first cues about the nature of institutional facts (or entities): as general social facts, they require collective intentionality, and in addition they require institutions for their existence. Up to now, we do not know neither what is an institution nor what kind of relation there is between institutional facts and institutions. To understand both, we need two investigate two further things: the assignment of functions by human agreement and the concept of status function.

5.2 The assignment of function

Another concept we need to consider is the assignment (or imposition) of function. Once again, Searle starts from a naturalistic observation. Humans possess a capacity they share with many other animals: the capacity to impose a function to a preexisting object, be it natural or artifact. Tools provide a clear example. Humans are notably adept as tool users, but other animals, such as beavers building dams and birds constructing nests, also exhibit this behaviour. The act of imposing a function indicates that functions are not OIFs. They do not arise directly from the physical features of the object in question. Instead, functions are always ORFs, since it is the agent that assigns them from the outside. But within the larger category of functions assigned specifically by human beings, Searle inserts a further distinction: the one between *agentive* and *non-agentive* functions. Agentive functions are functions that we assign to object for immediate practical purposes. With sentences such as “This is a rolling pin” or “This is

a kitchen faucet" we are marking «uses to which we put objects, functions that we do not discover, and that do not occur naturally, but that are assigned relative to the practical interests of conscious agents» (Searle 1995: 20). Those objects may be either naturally occurring ones (e.g. stones) or artifacts made specifically to perform that function (hammers and rolling pins). Non-agentive functions are instead assigned in the process of giving a theoretical account of some objects or phenomena under discussion. The example made by Searle for this subcategory is "The hearth functions to pump blood" in the context of giving a scientific account of how organism live and survive.

In both cases – agentive and non-agentive functions – we are dealing with ORFs, but the ascription of the former tracks the use to which *we intentionally* put the object. For this reason, objects that perform agentive functions cease to perform it whenever there is lack of continuous intentionality on the part of users. Not so for non-agentive functions, which are unaffected by collective intentions: hearts and livers keep functioning as hearts and livers despite our attention towards them.

With agentive functions we arrive, then, to the first interaction between collective intentionality and assignment of function: collective assignment of functions. As long as users collectively intend to act as if an object has a specific agentive function, that same object will perform its agentive function. A mass of wood and iron is a hammer as long as we collectively believe/think/act as if it is a hammer[39].

It is important to note that there is no need to state the function out loud or to make it salient every time in order to collectively assign it. Collective imposition of function is often unconscious or invisible, as Searle says. Money may have been evolved as a medium of exchange without people thinking "We are collectively assigning the function of medium of exchange to these metal coins" before any transaction. For example, people may have been born within a community that already use metal coins as medium of exchange and therefore adopt it as such, without questioning it or testing whether the assignment of function is in force. The very same act of using them as money demonstrates that the assignment is still in force.

[39] Searle does not specify what kind of collective action is to collectively intend something as such and such, but it is quite clear from the text that he intends something along the line of plain acceptance. That is, to collectively intend something as a rolling pin is to collectively act/behave as if that object is a rolling pin.

However the example of money takes already for granted the third element of our discussion: the concept of status function.

5.3 Status functions

We have said that in the realm of functions some of them are agentive, and that agentive functions require continuous intentionality on the part of users. In order to use something as a rolling pin, we need to continuously think of it as a rolling pin. However, the examples of agentive functions made so far, including the rolling pin case, are not so special in themselves. The reason is that objects performing such functions are also *physically suitable for them*. We regard a large log as a bench because the log possesses certain physical features, in terms of shape, toughness and stability, that enables it to perform the function of a bench. For the opposite reason, we do not usually think that sharp rock outcrops can function as benches. In both cases, outward physical features signal the suitability of the object in question in performing a specific agentive function.

Human beings, though, have developed the capacity to assign agentive functions also to objects (or people) which could not perform that same function solely in virtue of their physical structures. Sometimes we assign a function to an object or a person even in the absence of suitable physical features, but rather can perform the function only in virtue of the fact that there is a collective assignment of a certain *status*, and the object or person performs its function only in virtue of collective acceptance that the object or person has the requisite status. In these cases, a new specific type of function is assigned: the *status function*.

To illustrate the difference between 'ordinary' cases and status functions, Searle make use of the following parable. Imagine a tribe that constructs a wall around its collection of huts. Initially, this wall serves the function of keeping tribe members in and intruders out, relying on the members' tolerance to get over it. Now, imagine that over time, the wall decays until it becomes nothing more than a line of stones. Despite this decay, let's suppose that the tribe members continue to recognize the line of stones as a *boundary*. They understand that they are not supposed to cross it unless authorized to do so. In this scenario, the wall, which has diminished from a large physical structure to simply a line of stones, still serves the same function it did before. However, now it does so not because of its physical form but because the tribe members continue to

accept the line of stones as having a certain status – that of a boundary. They behave in a way that they consider appropriate for something they accept as a boundary.

Thus, the line of stones functions not because of its physical structure but because of the collective assignment of a status. With that status comes a function that can only be performed due to the collective acceptance of the object as having that status. We now have everything we need to talk about institutional facts.

5.4 The general structure of institutional facts

After presenting the last element, Searle claims that the assignments of status functions often take the form "X counts as Y in the context C". For example, a specific move in football is considered as scoring a touchdown, certain procedures are considered as electing a President of the Republic in Italy, and a particular position in chess is considered as checkmate. In each case, the term X identifies certain features of an object, person, or state of affairs, while the term Y assigns a special status to that entity or situation in specific contexts. As Searle tells us, the "counts as" locution names a feature of the assignment of a status to which function is attached by way of collective intentionality, where the status so attached and the accompanying function go beyond the sheer brute physical functions that can be assigned to physical objects (like ordinary agentive functions). Let's go back to the example of money.

When we accept that some metal coins or bits of paper count as money, we are applying the "counts as" formula for status functions, because satisfying the X term, "these metal coins" or "these bits of papers", is not by itself sufficient for being money, nor does the X term specify causal features that would be sufficient to impose the function of money on metal coins without collective agreement. This is the general procedure that gives place to institutional facts, of which we can now disclose a first, tempting definition. Institutional facts are facts that exist exists only in virtue of collective acceptance of something having a certain status, where that status carries functions that cannot be performed without the collective acceptance of that same status. Albeit not erroneous, as it is formulated this definition does not make evident the role of the word "institutional" in the expression. To understand that role we need to take a small step back and to focus again on the peculiarity of status functions.

We have just said that status functions are assigned independently of suitable physical characteristics on the part of the object. The lack of suitable physical features of the object in question is crucial for two reasons. First, it indicates that we are dealing with the imposition of a status function (and not just ordinary agentive functions). Secondly, it stress the importance of continued collective acceptance for the validity of assigned function. It is not enough that we agree with the original assignment that "This coin counts as money". Instead, it is necessary that the procedure of counting X as Y becomes regularized. When this happens, Searle says, the procedure becomes a rule, and rules of the form of "X counts as Y in the context C" are constitutive rules[40].

Constitutive rules, contrary to regulative ones, are rules that do not simply regulate an antecedently existing activity, but they create the very possibility of certain activities. To act according to them is somehow constitutive of that very act/behaviour. The canonical example of constitutive rules is chess rules:

> The rules are constitutive of chess in the sense that playing chess is constituted in part by acting in accord with the rules. If you don't follow at least a large subset of the rules, you are not playing chess. (Searle 1995: 28)

The idea of Searle is that institutional facts are possible only within a *system of rules* of this type:

> The systems of rules create the possibility of facts of this type; and specific instances of institutional facts such as the fact that I won at chess or the fact that Clinton is president are created by the application of specific rules, rules for checkmate or for electing and swearing in presidents. (ibidem)

The central point here is that to have institutional facts it is not sufficient to have the application of a single constitutive rule – that is, rules of the form "X counts as Y in the context C" – in isolation. It is necessary to have a system of these rules that combines and structures them. You cannot only have the rule that determines victory in chess without having at the same time the rules that dictate how chess pieces move through the chessboard: to play chess you must have both (plus probably others). The same goes for the institution of money and presidential election. They are possible only within a

[40] Actually, the status of the "counts as" formula is uncertain, since Searle later relegates it to «a useful mnemonic to remind us that institutional facts only exist because people are prepared to regard things or treat them as having a certain status and with that status a function that they cannot perform solely in virtue of their physical structure» (Smith and Searle 2003: 301).

whole system of constitutive rules. Institutions, for Searle, are precisely that: *systems of constitutive rules that enable the existence of institutional facts.*

Now we know the definition of both institutional facts and institutions. Institutional facts are facts that exist only in virtue of systems of constitutive rules, that is institutions. Constitutive rules are rules of the form "X counts as Y in the context C", where X is an object or a person and Y is a status function that is attributed to X not due to their physical features but uniquely in virtue of collective agreement. The same formula can be reiterated vertically to create hierarchical structures of institutional facts. Therefore, something that fills the Y-place in one context may fill the X-place in another one at a subsequent level of reiteration. For instance, when I produce certain sounds with my mouth, this action (X_1) qualifies as uttering Italian sentences (Y_1). Uttering sentences in such and such way ($Y_1 = X_2$), in turn, may qualify as making a promise (Y_2). In such and such context, making a promise ($Y_2 = X_3$) may constitute entering into a contract (Y_3). Entering into a contract ($Y_3 = X_4$) in such and such context may constitute getting married (Y_4), and so on.

Alternatively, there may be interlocking systems of iterated structures that operate at the same level. For instance, performing the speech act of promising (X) in front of a presiding official (C) counts as getting married (Y), but the same role of presiding official is the result of a previous imposition of status function that we have assigned to an individual in virtue of collective agreement. That is, the individual who plays the role of presiding official filled the X-place of another antecedent "counts as" formula, whereas the function of presiding official filled the Y-place.

Now we know what institutional facts and institutions are, as well as what kind of dependency relationship holds between them. The core of Searle's theory is already set out, but in order to apply it correctly to the case of news we need to add some further important details.

5.5 Deontic powers

After presenting the general structure of institutional facts, Searle invites us to consider the imposition of status functions as the conferment of some sort of *power* to the object or person that fills the X-place. Two clarifications are required.

First, power is for Searle the capacity to do something or to constrain someone else from doing something. Then the content of power status functions is always "S does A", where "S" is a placeholder for either a single individual or a group, and "A" stands for the name of an action or an activity, including refraining from doing something. Secondly, powers are always conferred to single individuals or group of people, even when the power comes from status functions that are imposed on objects (here's why the "S" in the above formula stands always for individuals). The reason is that the kind of power status functions bestow to the X-term is actually the power to give or deny rights, duties, responsibilities, expectations, obligations, permissions and alike. Searle call all of these "deontic powers". And deontic powers work always in relation to the possessor of the object, and not to the object itself. For instance, my five euro bill, which has the status function of money, gives me the right to buy goods up to the value of five euros; my Italian driving license gives me the right to drive cars and motorbikes with a certain engine size in the Italian territory, and so on.

Since the general structure of institutional facts suggests that we bestow power by means of collective agreement, Searle proposes the following formula as the basic structure for the conferment of deontic powers:

We accept (S has power (S does A))

Call it *the power-conferring formula*.

The simplest applications of the power-conferring formula are cases of enablements and requirements. In the former case we collectively grant power to some individual or group; in the latter we restrict their power. The case of money is a case of enablement that may be formalised as follows:

We accept (the bearer of X is enabled (the bearer of X buys with X up to the value of five euros))

Where S is "the bearer of X" and X is the five euro bill (requirements follow roughly the same construction).

The same structure of collective agreement works by interposition. So there may be power-conferring agency, such as the local Department of Motor Vehicles, that grants someone the right to drive in their own country by realising driving licences. In this case, Searle clarifies that we need to keep separated the intentionality of the agency

from that of the community at large, which makes the whole thing work in the first place. Here is the formula from the point of view of the surrounding society:

We accept (the agency creates (S is enabled (S drives a car)))

Examples like these are therefore cases of vertical reiteration of the power-conferring formula, in which we collectively agree to bestow power on someone, who in turn bestows power on someone else and so on, like matryoshka dolls. But it is fundamental to acknowledge that the starting point of every reiteration is always collective agreement of members of the surrounding society.

But what is the relationship between constitutive rules and the power-conferring formula? It is quite straightforward. We do not just accept that somebody has power, but we accept that they have power in virtue of some status. Searle (2005: 17) makes the example of US presidency. Fulfilling certain requirements makes someone president of the United States. This is an example of the X counts as Y in C formula, that is the constitutive rule. But, once we accept that someone is president of the United States, then we accept that he has the power to do certain things. They have the positive power to command the armed forces, appoints the White House staff, and to sign a bill from Congress into law within ten day of receipt, among others. They also have obligations, viz. negative powers. For instance they have to deliver a state of the union address. In this case we accept that S has power (S does A) because S = X, and we have already accepted that X counts as Y, and the Y status function carries with it the acknowledged deontic powers.

Searle actually believes that the power-conferring formula is what really sums up his whole theory on social reality, for he asserts that the exercise of deontic powers has supremacy over the objects to which the status function is imposed. The reason is that objects are primarily designed to serve agentive functions: governments and money are «just placeholders for patterns of *activities*» (Searle 1995: 57). The object that fills the X-place is only the standing possibility for the exercise of some deontic power. My five euro bill is a standing possibility of paying something or, more properly, a standing possibility to exercise my right to buy goods up to the value of five euros. The bits of paper in themselves have little interest for us.

5.6 Summary

Searle's framework on social reality helps us clarify what entities that depend on institutions for their existence are. Let's sum up.

Institutional facts (or entities) are facts that exist only in virtue of systems of constitutive rules, that is institutions. Constitutive rules are rules of the form "X counts as Y in the context C", where X is an object or a person and Y is a status function that is attributed to X not due to their physical features but uniquely in virtue of collective agreement. The assignment of a status function always comes along with the assignment of a deontic power of some sort (rights, obligations, responsibilities…) to the person/group that fills the X-place or the person/group who possesses the object which fills the X-place. The assignment of deontic power follows the power-conferring formula "We accept (S has the power (S does A))", which eventually is always grounded on collective agreement of individuals, even in cases of interposition. From a certain point of view, it could be said that the power-conferring formula is what really matters for the creation of institutional facts, for the exercise of deontic powers has a supremacy over the object or the person upon whom we impose the status function (the X-place). We are interested in creating functions because we matter exercising some power, regardless of the object or person at stake. Here's why sometimes institutional entities can be created out of thin air, without any object or person filling the X-place of the constitutive rule (corporation and virtual money are examples of that).

At this point, we finally have all the theoretical tools we need.

6. News as an institutional entity

My proposal is to treat news as institutional entities à la Searle which follow the subsequent application of the "counts as" formula:

> Information [X] counts as news [Y] when it is issued by a journalistic source [C]

This constitutive rule deserves an in-depth examination.

Information is what plays the role of the X-place in the formula. I am not committing myself to any specific account of information. Sticking to the actual debate on fake news, I will adopt a non-factive account of information. Therefore, a piece of information is just propositional content, and as an abstract object it can serve as the

basis for the imposition of a status function[41]. Notice that the presence of information in the X-place does not entail that news is a kind of information as in the case of conservative accounts, for constitutive rules are not definitions. Constitutive rules are just ways to indicate the supervenience of a property (a status function) onto a basis (the object) by means of collective agreement. They do not indicate an identity between the object and the status function.

What matters more here is the second part of the formula. Regarding the privative view, I regard news as a status, but here the notion of status draws on the specific Searlean framework. The assignment of the status of news is the imposition of a status function, because a piece of information acquires it independently from its own features. There is no physical feature in a piece of information that makes it automatically suitable for having the function of news: satisfying the X-term is not sufficient by itself for something to be news; just like being a bit of paper is not per se sufficient to be money. Instead, we collectively accept that a specific piece of information is news as long as that same piece is delivered by a journalistic source (the context C). Not any kind of information is news. A road sign indicating the maximum permitted speed on an urban road is a physical implementation of some well-formed and meaningful data. But we would hardly call it news: the characteristics of the context in which information is issued are important.

Incidentally, this is why conservatives fail to account for the difference between a piece of information reported by someone and a newscast report (above, p. 71). The reason they are different entities is that the former is an abstract object while the latter is an institutional fact. From a certain point view, the newscast report is still information and thus an abstract object. But from another perspective, it is an institutional fact because it is information issued by a journalistic source. But why does it make a difference? After all, context variation may seem a very small detail.

To explain why, we have to recall that there can be interlocking systems of institutional facts (§5.4). Uttering the ritual formulas for marriage in front of a presiding official is not the same as uttering them in front of a station manager: only the former

[41] There is an endless debate on what propositions are. For the purpose of our work, we settle for a simple and widely accepted conception of propositions as sharable objects of the attitudes (what we think, hope, desire) and primary bearers of truth and falsity, which are not determined in time and space and have no causal relationship with the knowing subject (McGrath & Frank 2023).

case counts as getting married. Features of the context are relevant, for sometimes the object or the person that fills the C-place had been in turn invested with a status function. This is precisely the case of presiding officials: people who have the deontic power of celebrating a marriage in virtue of collective agreement – a case of enablement from the power-conferring formula. But this is also the case of a journalistic source! Take the example of newspapers. Newspapers, whether paper-based or digital, are sources of information that have the deontic power of issuing news items or, alternatively, the power of letting people know what (information) is news. Like presiding officials, they too have been already through the imposition of another status function. The only difference is that with newspapers we are dealing with a case of vertical reiteration of the power-conferring formula.

In fact, the case of newspapers is more akin to the example of the Department of Motor Vehicles (above, p. 86). The Department of Motor Vehicles is an agency that has the power of issuing driving licenses. In the same way, *The New York Times* is an agency that has the power of issuing news items. But for both applies the principle of interposition for deontic powers: it is always collective agreement of the social community as a whole that bestow them on the agency in question. Following the example of the Department of Motor Vehicles, then, we may be tempted to accept the following initial scheme for newspapers:

We accept (N creates (I is enabled (?)))

Where N is the name of a newspaper and I is a piece of information (we can imagine any other journalistic source in the place of newspapers). What is missing from the above scheme is the kind of enablement that the piece of information is able to exercise. In fact, if "news" points to a status function there must be some sort of deontic power that is conferred to information by the status function of news.

To understand what deontic power news has, we need to take a step back and ask why we accept that newspapers have the power of issuing news items. To investigate this aspect is simply to unpack the matryoshka doll that characterises cases of vertical reiteration of the power-conferring formula. In fact, the reason why we confer, through collective agreement, the power of issuing news to newspapers is that newspapers are themselves the result of multiple applications of the "counts as" formula. In part, this was already mentioned earlier, when the example of the presiding official for marriages

was given: the C-place at a higher level can be a Y term from an earlier level. The same goes for newspapers in the constitutive rule that confers the status of news to information. Newspapers play the role of the Y term in an antecedent reiteration of the counts as formula, because *also newspapers refer to a status function*. Not all illustrated pieces of paper, or websites, which display articles are necessarily newspapers. For a collection of printed or digital articles to count as newspapers, it must be produced by one or more journalists. We may render this in the following constitutive rule:

A collection of articles [X] counts as newspaper [Y] if produced by journalists [C]

which is in turn a particular instantiation of the following rule:

A source of information [X] counts as journalistic source [Y] if produced by journalists [C]

But the object that acts as the C-term in those constitutive rules is again a status function: "journalist" indicates a status that someone gains in part through collective agreement.

It may sound odd to say that "journalist" designates a status, for when we talk about professions we are inclined to think that someone (the X term) gains the title in question (the Y term) solely in virtue of their physical or intellectual achievements. For example someone can be an attorney because they are law school graduates[42]. But this is only partially true. Collective agreement is still required, because the certification or authorization that confers the title of "attorney" needs to be collectively recognized and accepted as such. If the surrounding social community suddenly ceases to recognize the validity of the title that certifies the function of "attorney", that title itself becomes useless and people having it would not gain any deontic power from it. Having a lot of money does not give you the power of buying goods in a world where nobody values money as a medium of exchange. In the same way, to be a journalist is in part to be collectively recognized and accepted as such (often tacitly and unconsciously).

The fact that "journalist" indicates a status function is important in the economy of our discourse because it entails the possession of deontic powers. What are those

[42] Being graduated in law is not sufficient for gaining the title of "attorney"; in many countries you are also required to do an internship in a law firm (and sometimes even to pass an additional exam). The point is only to say why in the case of professions you still need collective imposition of status function in order to make use of that title.

powers? As with many other professions, there is no universal legislation that puts in writing the rights and duties of journalists. But we may resort to many influential resources, made by distinguished journalists, that point to the existence of a set of ideals that any journalistic venture has to accomplish as far as possible. In their bestseller *The Elements of Journalism* (2014), for instance, the two reporters Bill Kovach and Tom Rosenstiel characterize journalism as a dynamic social institution, and they pinpoint a set ideals or duties that every journalistic venture ought to observe:

1. Journalism's first obligation is to the truth.
2. Its first loyalty is to citizens.
3. Its essence is a discipline of verification.
4. Its practitioners must maintain an independence from those they cover.
5. It must serve as a monitor of power.
6. It must provide a forum for public criticism and compromise.
7. It must strive to make the significant interesting and relevant.
8. It must present the news in a way that is comprehensive and proportional.
9. Its practitioners have an obligation to exercise their personal conscience.

This group of nine principles sums up the duties – or in the lexicon of our power-conferring formula, the *requirements* – of journalists[43]. Journalists are required to aim at the mark of the truth, to be loyal to citizens, to verify information, to be independent from those they cover and so on.

The possession of the title of journalist carries with it these nine obligations that can be formalized in nine relative power-conferring formulae. For instance, we may say that "we accept (the journalist is required (the journalist is independent from those they cover))". In fact, these nine requirements are those that, as newspapers readers, *we expect to be observed.* Kovach and Rosenstiel are explicit on this point:

> [They] are principles that citizens *have a right to expect* [emphasis mine]. The principles have ebbed and flowed over time, but they have survived because they provide things that

[43] Actually there was an additional principle, the tenth, which claims that citizens too have rights and responsibilities when it comes to news consumption. I decided to drop it from the list because it did not pertain properly to the journalistic side (more on that below).

citizens need from the news in order to adjust to the demands of life in an increasingly complex world (Kovach and Rosenstiel 2014: 9).

Therefore, the assignment of the status function of journalist creates both requirements and enablements. On the part of the journalist, there is a set of obligations that coincides roughly with the nine principles taken from Kovach and Rosenstiel (2014). Again, they are in no way conclusive, and surely they are not fine-grained principles; but they are more than sufficient for our analysis. These obligations are not random limitations, but they are justified by the fact that the journalist has in return a powerful enablement: the right to decide what is news and what is not – or equivalently, what information is newsworthy and what is not. To be newsworthy, in a sense, is to be worthy of the public's attention. The journalist has the deontic power to decide what information is worthy of being given to the public, due to its potential relevance. They can decide whether to reject a report or to broadcast it on the assumption that it is, respectively, irrelevant or interesting[44]. That is why we said before that newspapers, like *The New York Times* or *The Guardian*, have the power to issue news items.

On the part of the news reader, instead, there is the creation of a set of enablements (or rights) that stems directly from the nine obligations. The reader is enabled (or has the right) to expect that journalistic products are worthy of their attention (newsworthy), truth-oriented, transparent with respect to sources, that they have been verified, and so on. And again, this can be translated to the relative power-conferring formulae. For instance, we may say that "We accept (the newspaper reader is enabled (the newspaper reader expects to read truth-oriented articles))"[45]. Furthermore, there may be a whole set of enablements for readers that depends on the right of reproaching, even harshly, journalists that do not live up to their obligations. For example, it seems normal to imagine a right to criticize journalists who publish unverified or even misleading information: "We accept (the newspaper reader is enabled (the newspaper reader reproaches journalists that publish unverified claims))".

[44] As a matter of fact, they may also make this choice according to ideological biases, as the huge literature on media bias in economics and sociology testifies (Gentzkow et al. 2015, Puglisi and Snyder 2015). But here we are talking about journalism from an ideal and purely normative perspective.

[45] Notice that the same right is not possessed by a reader of a satirical magazine. Readers of *The Onion* or *Private Eye* cannot expect to find truth-oriented articles when they browse through their websites. Again, a mere collection of articles does not by itself become a newspaper if it is not produced according to journalistic precepts.

This double perspective of journalist and readers is crucial for our inquiry on the deontic powers of news. In short, my thesis is that news inherits the deontic powers that come from the status of journalist. A news item has the power *to instil a set of expectations with regard to information in the news reader*. When we read a news item we expect that the article in question is such and such (newsworthy, truthful, scrupulously verified, orderly ecc.), and we expect that because news is the product of journalists. In the context of newspapers, we nurture a whole series of expectations that we would not nurture for cases of reported information. We do not expect our acquaintance to scrupulously verify the testimonial or perceptual sources from which she acquired the information that a terrible car accident happened nearby our home. On the contrary, we do expect that from the newscast! If we lately come to know that the reporter who covered the news behaved sloppy and did not bother herself too much in finding accurate information, we would feel disappointed and upset. We would feel entitled to reproach the reporter and the whole editorial staff that has vouched for the quality of her journalistic product. Whereas, normally, it seems ungenerous to treat in the same way an acquaintance of ours who happened to be passing in the vicinity of the accident. She has no duty of scrupulous verification or independence from the people involved in the fact.

We can finally return to the point from which the digression started and offer the complete version of the power-conferring formula for news items:

We accept (N creates (I is enabled (I instils the expectation in the reader of I that I is worthy of public attention, truth-oriented, scrupulously verified…)))

To understand this formula, it may be useful to use an analogy with other professions. Imagine that the driveshaft of my car gets broken and I need to repair it. I bring my malfunctioning car to a friend of mine, Sem, who is a mechanic and promises me to fix everything in a week. From the fact that Sem is a mechanic, a status function, I expect that Sem is skillful in fixing car problems. I therefore expect the cars Sem returns to his customers to be fully functional. The same applies to doctors that make a diagnosis, which we expect to be correct and precise, or attorneys who deliver a speech in the courtroom, which we expect to be convincing, and so on. The status attached to a profession creates expectations towards professional activity of the individual who possesses that same status – the activity being a performance or a material product. Why

should it be any different for the activity of journalists? Just as we expect a mechanic to return working cars, in the same way we expect journalist to deliver information that is newsworthy, true, verified etc. And just as I would feel entitled to reproach Sem if my car comes back into a worse condition, someone else would be entitled to criticize a journalist that have disseminated unverified information.

Everything said so far explains why news is an institutional entity. Institutional entities are those that exist only within a system of constitutive rules. Constitutive rules are rules of the form "X counts as Y in the context C", where X is an object or a person and Y is a status function that is attributed to X not due to their physical features but uniquely in virtue of collective agreement. Furthermore, the assignment of status function always comes along with the conferment of some deontic power in the form of enablements and/or requirements. As we have just seen, news is something that exists within a system of constitutive rules that develops vertically as well as horizontally: individuals who write articles according to journalistic practices counts as journalists[46]; collection of articles counts as newspapers when produced by journalists; information counts as news when it is issued by newspapers or any other journalistic source. These constitutive rules assign every time a different status function that holds solely in virtue of collective agreement: a collection of articles does not automatically play the role of newspapers if no one accepts that. The assignment of the status function of news confers to information many deontic powers that news inherits from the professional status of journalist. News has the power to enable readers to nurture a set of expectations that in normal cases of reported information we would not nurture. News consumers are entitled to expect that information read on newspapers is worthy of their attention, true, verified, independent and so on. Besides, they acquire the right to criticize and harshly reproach journalists that deliver information whose quality does not live up to their expectations.

[46] Notice the difference between someone *being* a journalist – having earned a certificate or other status indicators – and someone *acting* as a journalist. In my view, what truly matters in producing news is the latter. In fact, as we shall see (p. 130), fake news can sometimes originate from sources that are journalistic *in theory* but fail in practice when journalists do not rigorously adhere to journalistic standards. In such cases, they may hold the title of journalist but fail to *act* as one. Conversely, individuals who are not professional journalists can still produce genuine news if they follow journalistic norms – such as those working for student or community newspapers. I am grateful to Kate Scott for raising this important question.

In the next two sections, I argue against the privative view of news. I intend to point out that collective agreement plays a role in determining the value of news (§6.1) and no imitation of the appearance of the news is strictly possible (§6.2).

6.1 Collective agreement and the value of news

Fallis and Mathiesen have argued that the right history of production is what counts to confer the status of genuine news (above, p. 75). The reason is that having the right history of production guarantees the value of news, which is an epistemic value. Genuine news is worth because information encapsulated in it, so to speak, is verified, truth-oriented, generally error-free ecc. Genuine news items are like major fashion brands' products: they are valuable because they are 'refined'. On the contrary, Fallis and Mathiesen add, other counterfeitable objects do not get their value directly from their history of production. Money is one of them. Buyers and sellers collectively agree to assign a status function to bits of metal and paper; and for this very reason they are valuable. Fallis and Mathiesen, therefore, conclude that the value of genuine news does not depend on collective agreement, because we cannot confer epistemic quality to something through simple agreement.

I believe that Fallis and Mathiesen's stance is essentially correct, although incomplete. It is true that (genuine) news has an epistemic value that comes from the observance of the rules of journalism, but I don't think that news is valuable only for this reason. To be precise, I suspect that, in their argument, Fallis and Mathiesen are surreptitiously using two different meanings of "value", one technical and one not. When they talk about money and say that US dollars have value due to collective agreement, they are using the word "value" with the technical meaning of status function. Bits of paper called "dollar" have value in the sense that they have acquired a status function through collective agreement that confer them some additional properties (e.g. being legal tender) that they could not have solely in virtue of their initial features. I think there are no doubts on this, for they explicitly mention Searle's theory to account for the case of US dollars (Fallis and Mathiesen 2019: 9).

On the other hand, though, when they talk about genuine news their focus shift to a non-technical notion of value. To have value, now, is to be the outcome of a high-quality production process. In fact, they suggest to compare fake news to counterfeit

versions of brands that are valued for their quality[47]. However, if you use two different meanings of value you cannot safely conclude, as Fallis and Mathiesen do, that money and news are not comparable in terms of value. I believe that what they say is true: news with the right history of production is (usually) more epistemically valuable than news that did not undergo the correct process of production. But what if, in addition to epistemic value, news had value for the same reason that money has value?

Take again the example of money. It is true that in the case of pure fiat money, like euros or current US dollars, you cannot find any intrinsic value other than the one bestowed by collective agreement[48]. The five euro bill has value because it allows you to buy goods and services that together are worth a maximum of five euros. This is a deontic power – the right to buy up to a certain value – that is conferred following the power-conferring formula, which in turn is grounded in collective agreement ("We accept that…"). But in §6 I tried to show that something analogous happens with news.

The imposition of the status of news is governed by its proper constitutive rule: information counts as news when published on newspapers (or any other journalistic source). Constitutive rules dictate how specific status functions are bestowed on people or objects trough collective agreement. Therefore a piece of information is news *because we collectively accept that information issued by newspapers/newscast/news bulletin/ ecc. counts as news*. At the same time, since the imposition of a status function comes with the conferment of deontic powers, we also give news some deontic power. News has deontic powers inasmuch as it can instil a set of expectations in the news readers that would not be elicited in the same way by any simple piece of information. This enablement stems from the fact that pieces of information which fill the X-place in the constitutive rule for news are issued by journalistic sources, which are in turn institutional entities. Newspapers, for instance, are institutional entities because they

[47] Perhaps there is a further meaning of "value" in Fallis and Mathiesen's analysis; a sense that can be found in other privative accounts too. According to this third meaning, genuine news is valuable because it can efficiently satisfy an epistemic need. Here the notion of value is still a technical one, but it is more akin to the notion of value that pertains to economics: goods have value because they are able to satisfy needs or wants of consumers. Value is here a synonym for economic utility.

[48] Beware of the fact that here we are taking for granted the correctness of Searle's theory for money as institutional fact. I am not claiming that Searle's interpretation of money is also correct from the standpoint of monetary theory. In fact, the general accepted view of monetary theorists on fiat money is that unredeemable bits of metal and paper have value because the state, or the sovereign, backs them with its power (see Dowd 2001 on this).

exist only as products of journalists; a mere collection of articles written by random scribblers is not ipso facto a newspaper.

And as we go down the pile of status functions imposed one on the top of the other, we find that the base of this structure coincides with the status function of journalist. News inherits its deontic powers through a chain of imposition of status functions that starts from the status of journalist. The journalist is the one that has duties of verification, commitment to the interests of the readers, independency from the subjects covered… and therefore we expect her products – i.e. the news – to live up to those duties. But if the status of journalist, as any other profession, is partly granted by collective acceptance of the possession of that status, then we find once again that at the origin of all this there is always some form of collective behaviour. The difference between money and news is only that, in the former case, the role of collective agreement is immediately manifest. You do not have to go very far down the pile of status functions to get to its base. In the case of news, instead, you have to dig a little bit deeper, because news is the end product of a longer chain of impositions of status functions that starts with the profession of journalist.

The moral of the story is that news has value in the same way money does: through collective agreement. This in turn does not mean that news cannot have any other form of value; in fact, it is true that it also has epistemic value, like Fallis and Mathiesen rightly observe.

6.2 On imitation of news appearance

The second objection that I intend to make to the privative view concerns one of its core claims: the imitation of news appearance. Privativists in fact argue that non-genuine news – i.e. fake news – imitates superficial aspect of genuine news in order to pass itself off as authentic (cf. chapter 1, §7). This move would enable fake news to represent itself as something that possess the status of news, even if it does not really have it, because it does not have the right history of production. In short, fake news is information that looks like news but in fact it is not. What is wrong with this analysis? Two things strike me as unconvincing, and each of them deserves unpacking.

Firstly, the privative view seems to take for granted that a universal set of superficial features characteristic of news does exist. In fact, in order to say that fake news imitates

the appearance of news, a defender of the privative view must also presuppose that there is a group of acknowledged outward features of genuine news that can be imitated to begin with. These features must behave as recognisable traits that signal, by themselves, the possession of the status of news. The existence of such traits meets the necessity of signalling the possession of a status, which is something that cannot be tested physically. Searle (1995) too has reflected upon this matter. He in fact claims that a reliable test for the presence of (the status of) institutional facts is whether or not we can codify their constitutive rules explicitly. In fact many of them – including money, marriage, private property etc. – have been extensively codified in the explicit laws of many modern societies. But why codification matters?

The reason is that there are a number of social facts concerning which we are uncertain to whether assign or not to them the label of institutional facts because they lack codification but, at the same time, show typical features of institutional facts. The examples made by Searle here are the ones of friendship, cocktail parties and dates. Up to now, and as far as we know, modern societies do not have codifications that establish whether or not a social fact is a friendship. The ascription of the status function of friendship is made on a case-by-case basis, but it is clear that people harbour certain sorts of expectations towards a friendship which they do not have from an identical set of arrangements about which they do not believe that it is a friendship, Searle says. If at some point in human history the assignment of the status function of friendship will be more pressing, we can imagine this informal institution becoming explicitly codified. We will know that a friendship arises every time such and such conditions are fulfilled.

A second benefit of codification is that it makes it easier to recognise an institutional fact in force. We have said that institutional facts arise from imposition of status function out of collective agreement. But we also know that status functions are by nature assigned to something or someone regardless of their physical features. This implies that institutional facts cannot be read off from brute physical features. That two people are married cannot be inferred by, say, the fact that they always spend time together or that they live together in the same place. In the same way, we cannot say that someone is entitled to drive in the Italian territory only because we saw them driving a car within Italian boundaries. This is the same problem raised before, talking about the

status function of news and the need of having recognisable traits. So how do we actually know those institutional facts?

The process of codification usually solves this problem by setting *status indicators*. Status indicators are reliable cues for status functions. When authentic and truthful, they indicate to someone the presence of the relative institutional fact. For instance, you can deduce that someone is married by checking whether they wear a wedding ring. Or you can verify that someone is entitled to drive in Italy if they possess a valid driving license. Wedding rings, driving licences, police uniforms as well as plenty of other things that signal a specific status are all status indicators. Status indicators are useful because they signal the maintenance of institutional facts. For example, by signing a document you may create an institutional fact (e.g. a new contract is now in force), but the written signature is what indicates the continued existence of that institutional fact through time. Status indicators do not constitute institutional facts – that is only possible through imposition of status functions out of collective agreement. They instead enable us to identify pre-existing institutional facts.

Recognisable traits of news are therefore nothing more than status indicators for the presence of the institutional entity called "news". When privativists say that fake news imitates genuine news appearance they are simply saying that fake news deceptively adopts the status indicators of genuine news. But the question now becomes "Do status indicators of news really exist?". We have just said that status indicators are fixed by codification. The reverse side of 1 euro coins, the common one among European countries that adopts euros as legal currency, must display a map of Europe: this is a recognisable trait, that is a status indicator, for money in the Eurozone. The design of this coin has been decided by the European Commission in collaboration with the European Central Bank. The fact that 1 euro coins must present a map of Europe in their reverse side is then explicitly codified. But what about news?

As far as I know there is no explicit codification for news. If so, there are not explicit rules that settle the status indicators for news. Contrary to money and other institutional facts, no one has established once and for all what has to be the official appearance of news. No central authority has dictated that a news article must present such and such

features – also because such authority does not exist for journalism[49]. Sure, there are always evolving conventions among practitioners that regulate this aspect, but they are in no way definitive or universally accepted. It seems, then, that news is an institutional entity more akin to friendships, dates and cocktail parties. Notice that absence of status indicators should not be confused with absence of reliable cues for news trustworthiness. They are two distinct aspects of news.

Recently, some empirical evidence has pointed to the fact that the appearance of online news sources might contribute to augment or diminish their perceived credibility. Wobbrock and colleagues (2021) have conducted two studies that solicited participants' credibility perception by isolating article presentation from article content. Article presentation includes style of the headline, number of images, font size and many other details. The central finding of the overall experiment is that certain presentational factors, either in isolation or in combination, do affect the perceived credibility of online content[50]. Furthermore, tested participants were often able to distinguish articles coming from credible news sources from articles coming from non-credible news sources, all by just sticking to outward features. But this is hardly a proof for the existence of status indicators for news.

In fact, although there might be a combination of factors that increase or decrease news credibility, this observation misses the point. We were not investigating cues for credibility. We were looking for indicators of the presence of the institutional fact of news, regardless of its subsequent credibility. In other words, we were interested in knowing whether there are outward features of information that automatically signal the presence of the status function of news. And it seems that they are not, so far.

At this point someone may reply that it is not that status indicators for news do not exist; it is just that we are looking for them in the wrong place. In fact, even though

[49] Legislation on journalism varies significantly across countries. The Italian *Association of Journalists*, for example, establishes the necessary procedures for becoming a journalist in Italy and may sanction journalists who do not observe their code of conduct, all the way up to removing them from the register of journalists in extreme cases (although it rarely does and the Association is progressively losing its prestige, and so its power). But things work differently elsewhere. In Germany and Spain, for instance, the legal status of "journalist" granted by inclusion in a dedicated register does not really exist, but there are laws that regulate the nature and the activity of journalists that refers to the concept of autonomy, deontology and access to the journalistic profession.

[50] Specifically, the authors of the study have found that a certain balance among factors (e.g. not to may images, medium size font ecc.) elicits the greatest amount of perceived credibility. They called this balance "a Goldilocks zone", in honour of the protagonist of the children's tale *Goldilocks and the Three Bears*, whose protagonist, Goldilocks, looks for porridge that is neither too hot nor too cold.

there might not be stylistic features of news that are universally accepted, there are nevertheless recurrent structural features of news that may deserve the role of recognisable traits. For example, a news item is always composed of a title (or headline) and a text. Therefore the double presence of those items must be a reliable cue for the presence of news. Once again, though, I believe that this further reply does not hold water. It confuses stereotypical manifestations of news, such as news articles, with news in general. A quick look to the history of news shows that a news item was not always characterized by the double presence of a title and a text. Between the 19th and the 20th century, journalists have massively deployed telegrams as a means of communication (Charnley 1966, Hampton 2004). The content of a telegram could include news updates, urgent information, or directives. While some telegrams might feature a brief heading or title to indicate the subject, many did not. The main priority was clarity and brevity to ensure the message was quickly understood. And it does not seem valid to claim that telegrams managed by journalists, but devoid of titles, were not news items for all intents and purposes. The same goes for the newscast: are we willing to say that information given through newscast is not news because it physically cannot have a title and a text? The reality is that the structure of news changes through time because it depends in part on technological innovations. That is why online news does not have the same features of print news, which in turn is different from news in the form of telegrams, and so on.

Actually, news shares this peculiarity with many other institutional entities that have not been codified. Uncodified institutional entities may present different manifestations through time and across cultures, because they do not pay the price that comes with explicit codification. As Searle says, when you set the correct procedures for the creation of an institutional fact you also lose «the flexibility, spontaneity, and informality that the practice has in its uncodified form» (Searle 1995: 88). This passage is of paramount importance for us. Whenever the conditions for the creation of an institutional fact have not been already codified, we may feel freer to assign the related status function on a case-by-case basis, for the decision to assign it is left to the opinion of people involved in the situation. So, for instance, two people may consider an encounter as a date because they judge on the basis of their opinions that this is the case; but two other people, if confronted with the same case, may not. At some point in the

future, the society as a whole may decide that there are some necessary and sufficient conditions for dating that should apply universally. These conditions dictate what kind of properties the object that fills the X-place – here, the encounter – should possess in order to gain the proper status functions – the one of date. Once those conditions are explicitly codified we are sure that if they occur we are dating someone. But at the same time, we are no longer free to decide whether an encounter is a date or not. Borderline cases that seem dates to us may not count as dates anymore and vice versa. Codification inevitably comes with rigidity.

And what happens with dates also happens with news. Since there are no universally accepted status indicators for news, people will decide whether to believe that a piece of information is news on a case-by-case basis. Specifically, this judgment will hinge on *whether the reader believes that the piece of information is issued by a journalistic source or not*. If so, any piece of information provided by that source will be treated as news by the reader, as long as she has no reason to think otherwise (e.g. she finds out at some point in the future that the source does not abide by journalistic standards). Then, contrary to what privative view seems to maintain, my hypothesis is that readers do not ask themselves "Is this (genuine) news?" any time they stumble across a piece of information. They do not recognize news by relying on the verisimilitude of its appearance, since there are no status indicators for news, but from the very fact that they found it on (what they think is) a journalistic source. If they trust the source – that is, if they believe that the source really is a journalistic one – readers will behave as if any piece of information provided by that source is news, *regardless of its appearance*.

At the same time, however, this does not mean that the appearance of news is irrelevant. As we will see also in the next chapter, news consumers have expectations regarding how news is published or presented to the public. These expectations partly stem from certain journalistic conventions (e.g., the 5 Ws model), which create a form of 'social codification' that is unofficial but widely accepted and recognized. Such conventions evolve over time, reflecting cultural, technological, and stylistic changes, making them flexible but, at the same time, less universal compared to the formal codifications required to attain a true status indicator. And here comes the flexibility of news as uncodified institutional entity: to a certain degree, readers may accommodate also deviant outward features of information if they still think that it is provided by a

journalistic source[51]. Therefore no perfect imitation of news appearance is necessary to pass something off as news. In some cases, it might be sufficient that the reader thinks that the piece of information in question is issued by a journalistic source.

To sum up, the privative view of fake new maintains that fake news imitates the appearance of genuine news to appear authentic, despite lacking proper production history. The argument presumes that there are identifiable, universal features of genuine news that can be imitated. Using the language of Searle's theory, this means that status indicators of news do exist. But in order to have status indicators, an institutional fact must be explicitly codified. However, unlike money or marriage, news lacks such codification. Status indicators, like wedding rings for marriage, help identify institutional facts, but no explicit indicators exist for news, because constitutive rules for news has never been explicitly codified. News is more like friendship or dates, which are not explicitly codified, and sometimes they are recognized based on context or individual judgments. Regarding news, those judgments will depend mainly on whether the reader will treat the source of emission of information as a journalistic source. If so, she will believe that what is reading is news, independently of the superficial features that the piece of information will exhibit. Thus, to pass information off as news there is no need to imitate news appearance.

Conclusion

In this chapter, I have asserted that the problem of the conceptual necessity for the conservative view of fake news stems from an excessively minimalist conception of what news is. News is not simply a type of information, but indicates a status: a

[51] When I specify that this happens only to a certain degree, I mean to suggest that things may be far more complex in real cases. In fact, when the news consumer has to decide whether the source in front of them is a journalistic source or not, they are certainly influenced by its appearance as well, in a sort of feedback cycle. So people will adjust their expectations only within a certain range. Where this range is situated is once again a subjective matter, but there is room for speculation. For instance, if an article on a newspaper is rife with grammatical errors, presents incorrect references to places and people, and does not exhibit ordinary credentials of a newspaper (like its logo), it may be very hard for a reader to convince themselves that they are reading a genuine journalistic product. It is unlikely that they will adjust their expectations about news in order to include that case. Again, this process cannot be measured objectively, because there are many subjective elements. We do not know the exact expectations that a news consumer harbours on journalistic products, and if their idea of what it means to deviate too much from a norm. My aim here, though, is not to make an empirical investigation, but only to hypothesize how we approach to news items. As we shall see in the next chapter, fake news lurks into the folds of this phenomenon. If our expectations on journalistic products are flexible, and if the aforementioned feedback cycle holds, then *in some cases* what is fake news for me may not be fake news for someone else.

property that supervenes on information when it meets certain standards. The privative view of fake news supports this vision, but ends up using a notion of status that is imprecise and not framed at a philosophical level. Moreover, the privative view commits to the idea that fake news imitates the appearance of real news, but this thesis is unrealistic. There is no universal checklist of the outward characteristics of real news that can be used as a model for imitation. I hypothesize instead that our expectations, as readers, of what constitutes news are flexible: if we believe that certain information is published by a journalistic source, then we are willing to consider that information as news and accordingly accommodate its outward characteristics. My thesis, therefore, is that certain information counts as news if it is published by a journalistic source. I draw this idea from the framework developed by Searle in *The Construction of Social Reality* (1995), and treat news as particular institutional entities. News is considered an institutional entity because its existence depends on the existence of a certain institution – the journalistic source – and because news has the deontic power to instil certain expectations in a reader; expectations that are not normally elicited by 'mere' information. Among these is the expectation that the information is worthy of public attention, truth-oriented, verified, and so on.

In the next chapter, we will define fake news in a way that is partly related to the conception of news as institutional entities. We will then see how expectations play an important role in determining our judgment of what is fake news and what is not. The main thesis that I will try to support is that fake news is news that somehow does not live up to our expectations about news.

Chapter 3

Fake news: an institutional account

1. The definition of fake news

In the previous chapter I contended that news is a status function that is conferred to a piece of information according to the following constitutive rule:

Information [X] counts as news [Y] when it is issued by a journalistic source [C]

where the C-term is in turn a Y-term for another constitutive rule that makes the status of journalistic source dependent on the status of journalist:

A source of information [X] counts as journalistic source [Y] if produced by journalists [C]

In light of this theory of what news is, I propose to define fake news (FN) in the following way:

FN: fake news is news that does not deserve its status.

To better understand this definition, it is convenient to split FN in two complementary interpretations: FN_A and FN_B.

FN_A: fake news is news that significantly violates norms of journalism.

FN_B: fake news is news that does not live up to a sufficient degree to realistic expectations towards the products of a journalistic source (journalistic products).

Both FN_A and FN_B describe the same phenomenon, that is FN, but from different perspectives. FN_A stresses the normative breach brought about by fake news, while FN_B highlights the relation between fake news and expectations of the news consumer. However, in both cases we are dealing with news that does not deserve its status, either because it violates norms of journalism or it does not live up to realistic expectations towards journalistic products; and as we shall see, these two alternatives are in fact interrelated.

Keeping these two sides distinguished may help to get important aspects of the phenomenon of fake news that FN, which is more abridged, may not get across alone. The aim of the current chapter is to unpack the meaning and the implications of FN_A and FN_B. To do it, I first need to answer the following questions:

1. What are those realistic expectations at stake in the definition of fake news? And what does it mean that a piece of news does not live up to them?
2. What are norms of journalism and how can they be violated?
3. In what sense fake news does not live up to our expectations "to a sufficient degree"? Or equivalently, in what sense it "significantly" violates norms of journalism?
4. Is fake news still news?

Section §2 is fully dedicated to questions 1, 2, and 3. There I elaborate more on the norms of journalism, describing what they are and how they are related to our expectations towards journalism. Then, in §3, I reply to 4 with personal considerations. Mine is a defence of the conservative view of fake news, albeit from a completely different premise. In fact, the starting assumption is no more that news is information of some kind, but that news is the institutional entity that I have presented in the previous chapter. Finally, in §4 I focus more on the pragmatics on fake news, with the goal of identifying interesting cases of fake news that fall under FN.

2. Norms of journalism and realistic expectations

In the previous chapter, I said that news has deontic powers. Specifically, news has the power to instil a set of expectations in the news consumer. Whenever we are about to consume it, we expect a news item to be worthy of our attention and interest, truth-oriented, verified in its content and so on and so forth. This deontic power is inherited directly from the profession of journalism: the journalist has their own set of rights and obligations that stems from their own status function (i.e. the profession itself). The main deontic power of the journalist is to decide what information is newsworthy, that is what information is worthy of the attention of the public. But to exercise it, the journalist is expected to observe a strict list of obligations and duties – the requirements.

I have already presented some of those requirements in the second chapter (§6), drawing on Kovach and Rosenstiel (2014). But what do those requirements represent? Kovach and Rosenstiel depict them as ideals that every journalistic venture should try to accomplish. For the sake of our discussion, let's call them *norms of journalism*. The norms of journalism, like the journalistic standards of the privative view, stand for all those kinds of requirements that are part and parcel of the profession of journalism.

They are the rules that define the way an ideal journalist should work[52]. These requirements, we said, have also an epistemic import. Some of them might also be seen as good epistemic practices in disguise. For instance, the norm that requires a scrupulous verification of information may resemble a simple epistemic practice of collecting sufficient evidence before forming a belief or issuing a judgement; the norm that requires independence from sources may be the analogous of an epistemic requirement of unbiasedness; and so on.

Norms of journalism, though, do not represent only ideals and rules of the profession of journalism. From the perspective of the news consumer, norms of journalism represent the *expectations* that they have for a journalistic product. These expectations are realistic when the news consumer knows sufficiently well the norms that the journalist should follow, so that the former can notice whenever the latter breach those norms thereby failing their duties. In all other cases, when the consumer is not familiar with the norms or has a distorted view of them, their expectations towards a journalistic product may not be realistic. Whenever we listen to the news bulletin, we can realistically expect that the reporter who is delivering us information has verified it beforehand, that this piece of information might be worthy of our attention, and so on. But if we expect the journalist presenting the news bulletin to take a position on the presidential election and tell us their orientation, or suddenly start using the vernacular to comment on the news; then probably our expectations around the journalistic product – in this case, the news bulletin – are not very realistic (even though, as we already said, there might be differences in expectations regarding norms of journalism due to contextual or cultural variations).

Whether realistic or not, expectations instilled by news together with norm of journalism are two sides of the same coin, the coin being deontic powers of journalists. The former is deontic powers from the perspective of those who passively consume the products of journalism, while the latter is deontic powers seen through the lens of the profession itself. The two component of the disjunction in FN, FN_A and FN_B, are grounded in these two sides of news: norms of journalism and their related expectations. FN_A is a definition of fake news from the perspective of the news producer, who can

[52] My account is not the first to highlight the role of journalistic norms in the definition of fake news. The one put forth by Pepp, Michaelson and Sterken (2019) already addresses the potential relation between fake news and journalistic standards.

either observe or violate the norms of journalism. On the other side, FN_B is fake news from the perspective of the news consumer, merely understood as the addressee of news. Note, however, that real news consumers judge if something is fake news or not on the basis of their personal set of expectations, which are not always realistic. So FN_B takes the perspective of an idealized subject (more on this below). Since ascertained violations of the norms of journalism lead to disappointment of the realistic expectations that we have about journalistic products, then FN_A and FN_B actually describe the same phenomenon: news that does not deserve its status.

There are additional considerations to be made about FN_A and FN_B. First of all, the notion of realistic expectations in FN_B does not always reflect reality. In fact, real news consumers have their own, subjective expectations of what is a journalistic product, and those are not necessarily realistic. Not everyone harbours the same expectations towards the activity of journalists, there is both individual variability and context specificity. The former means that different news consumers have different expectations simply because they have discordant opinions about the duties of journalists. Sometimes they do not know all the journalist's obligations, sometimes they have a distorted view of them. Context specificity, instead, concerns the fact that obligations of journalists may vary a bit across places due to cultural variation. Europeans have different ideas than Americans about the role and the function of journalism, for instance; and Northern Europeans think differently from Southern Europeans on the same issue (Hallin & Mancini 2004). Nevertheless FN_B assumes that the expectations of news consumers are realistic, that is grounded in valid opinions of what are the norms of journalism that have to be observed. So it takes for granted that readers always know what it takes to deserve the status function of news.

Another aspect to mention is that the relationship between FN_A and FN_B should not be considered a relationship between contingent facts. What I mean by this is that someone may be tempted to wonder what happens when a news item actually violates the norms of journalism without a reader realising it, perhaps because they do not have sufficient evidence of a violation. Below you find a useful example of this situation taken from Anderau (2021). Here is the slightly refined version of his case:

(1) Cancer-causing soft drink

> Lois writes a story claiming that there is a link between drinking soft drink X and cancer, citing well-respected scientific sources. When the story is published, readers have no reason to believe she violated norms of journalism; in fact, they believe it is genuine news. One year after the publication, it is revealed that the scientific evidence Lois has based her article upon is false and that the scientists publishing the story were in fact working for the producer of a soft drink named Y, a competitor of X. Two years after the publication, it is also revealed that Lois had been bribed by company Y in a plan to undermine competitors such as X and that she knew the science she was quoting was fraudulent.

Here is how the reasoning might sound. Lois's article is not fake news, according to FN_B, for the news consumer has no available evidence of a violation of norms of journalism. For them, the article meets the realistic expectations towards a journalistic product. One year after the publication the assessment is the same: even though the scientific proofs supporting the link between drink X and cancer turned out to be false and biased, readers cannot say for sure that the article published by Lois is fake news. Maybe she is a diligent and renowned reporter, and despite having published something false she may have complied with the norms. But two years after publication things are different: now readers undoubtedly know that Lois did not observe the norms of journalism and she did not fulfil her obligations as a journalist. The case in CANCER-CAUSING SOFT DRINK is now fake news, according to FN_B. However, even though the evaluation of news consumers changed over time as new evidence became available, Lois's article was fake news since the beginning. That is because her article was already published in full violation of some important norms of journalism. The violation was already there, but it went unnoticed for two years. Therefore, according to FN_A Lois's article was fake news in all three scenarios. If so, we have a time frame in which FN_A and FN_B predict two opposite outcomes. Therefore, FN_A and FN_B do not capture the same phenomenon.

However I think that this objection is flawed at its origin. First of all, as I anticipated in the first chapter, my definition of fake news is primarily ameliorative, and it does not aim at capturing the concept of fake news in terms of necessary and sufficient conditions as in the case of a descriptive project. Besides, the whole argument in my advice mistakes FN_B, which is a definition of fake news, for a heuristic that should help us identifying fake news. It gets wrong the role of FN_B. FN_B actually states what is fake

news regardless of whether real news consumers recognize it in concrete situations. It tells the kind of relationship that fake news entertains with our realistic expectations of what a journalistic product should be; so it works always *under the assumption* that the news consumer has evidence of a violation. To further understand why the objection against FN_B does not hold up, consider the standard definition of lying: to make a believed-false statement to another person with the intention that the other person believe that statement to be true (Mahon 2016). Criticizing FN_B on the basis of a case such as CANCER-CAUSING SOFT DRINK is like saying that the standard definition of lying is problematic because there are many cases where it cannot be determined with certainty whether an alleged liar has the intention to deceive or not. This would mean confusing definitions with criteria for identification; that is what I meant when I said the argument misunderstands the role of FN_B. The aim of definitions, whether they are descriptive or ameliorative, is to say what the definiendum is or should be, not to help us navigate the actual cases; the latter is at best a possible employment of a definition, not its primary goal.

Finally, it may be argued that FN_A and FN_B do not capture the same phenomenon, and so they cannot be subsumed under one unique definition (FN). But the extended discourse on FN_A and FN_B was aimed precisely at avoiding an issue of theoretical unity. In fact, we saw how violations of the norms of journalism and disappointment of realistic expectations towards a journalistic product are intertwined. Realistic expectations towards a journalistic product are grounded in observance of norms of journalism, so that violations of the latter are linked to disappointment of the former and vice versa. Therefore, the two components of FN do not capture two separate phenomena; it is the same one from different points of view.

2.1 Typology of norms

Up to now I have described norms of journalism in a cursory manner. Some of the norms cited so far are the ones listed by Kovach and Rosenstiel (2014). Pepp, Michaelson, and Sterken (2019: 75-76) have reformulated this list by extracting its salient points: commitment to the interest of citizens (or other members of a society); careful verification of putative information received; active seeking out of information that is or will be important for members of society; transparency about sources and

conflicts of interest; independence from those covered and from sources; proportionality in the amount of coverage given to events relative to their likely significance; and historical and social contextualization and synthesis of information.

What immediately stands out from this list of norms is the fact that they all operate upstream in the news cycle. They describe a set of commitments that regulate the practices concerning the collection, analysis and processing of information before it acquires news status. We may call them norms of *news production*. Roughly, these norms are comparable to a list of good epistemic practices for journalists. They do not guarantee any degree of success; scrupulous verification may nevertheless lead reporters to a failure for other self-independent causes (e.g. their usually-reliable source was wrong). But they are valuable in themselves and, most importantly, they represent a first round of check and filters that set news apart from simple information delivered by testimony or other individual sources.

However, norms that characterize the journalistic enterprise does not boil down to norms of news production alone. I maintain that there is also a whole range of rules that should regulate how information is finally presented to the reader. This set consists of what I call norms of *news publication*. Regarding standards of news production, norms of news publication ought to represent our background expectations for journalistic writing too, albeit at different level. They operate downstream in the news cycle, because they come into action after information has been gathered and processed according to norms of news production. But what are norms of news publication? Where do they get their theoretical justification?

We might initially consider the well-known conversational maxims put forth by Grice (1975) as a starting theoretical model for norms of news publication. My research, in fact, is very much focused on the expectations at play in the linguistic practices of journalism. As we see in Grice's work, the existence of certain expectations reveals that there is a specific type of norm in place that underpins those very expectations. Similarly, a journalistic product is usually expected to be adequately informative, true, relevant, perspicuous, unambiguous, brief and orderly as much as possible. Therefore, it is tempting to take Gricean maxims, or a substitute thereof, as the archetype for norms of news publication. This choice is palatable also because Gricean maxims (or a derivative thereof) would provide norms of publication with a solid theoretical

justification and an established tradition in the literature. Nevertheless, I believe that norms of news publication are not directly reducible to Gricean maxims. The reasons that drive my assessment can be summarized in the following four points:

- Violations of Gricean maxims can be strategic, not so for norms of journalism;
- Journalistic communication does not take hold in a context of mutual adjustment;
- Norms of publication are stricter in their application and allow for a lower degree of permissibility of violation (at least in certain circumstances);
- Trust in the journalistic source gives priority to observance of norms of publication;
- The style is of great importance in journalistic writing.[53]

The first thing to notice is that Gricean maxims were not intended as a style guide. In many cases, a violation of a Gricean maxims is only an apparent violation to generate an implicature in various ways. In this sense, violations of Gricean maxims can be strategic actions. For example, because two maxims clash and it would be impossible to abide by both, or as with flouting, where they are not really violated (i.e. you may 'say' something untrue, but you do so to implicate something true – so truth is still the overall intention). Instead, violations of journalistic norms do not function in this way: a journalist cannot, say, report falsehoods with the intention of generating a true implicature for the audience. The reason is that norms of journalism are not the same sort of thing as a maxim. They are a style guide. The maxims are not. Hence, each time we will talk about "violations" of the norms of journalism we are considering only manifest, non-strategic infringements of stylistic guidelines.

Another reason of diversity is that, contrary to talk exchanges, information delivery accomplished by journalism does not take hold in a context of mutual adjustment. Whenever we engage with a face-to-face conversation we usually have the possibility to challenge the speaker's statements, to ask for clarification, and in general to make any conversational move that can pile on the others. This opportunity for interaction gives

[53] These points are not exhaustive, and other compelling reasons may be put forward to differentiate the rules at play in ordinary conversations from those at stake in journalism. Mine is just an attempt to sketch journalistic communication as something that deserves its own independent characterization (and its own rules as well).

place to a context of mutual adjustment: the speaker can modify their statements in response to the requests and the demands of the addressee. Challenges like “What do you mean with that?” or “How do you know this?” are prompts for adjustment, that in normal circumstances should drive the speaker to better formulate their statements.

However, the kind of communication that unfolds between the journalist and the news consumer is different. The news consumer cannot ask for clarification or challenge the journalist’s statements, because their communication is one-way: the news consumer and the journalist are physically distant and/or the former is getting to know the statements of the latter through a written text. In any case, the result is that both are not immersed in a context of mutual adjustment when they enter into communication.

This discrepancy generates different kinds of commitments. In fact, both the journalist and the news consumer are aware of the absence of a context of mutual adjustment when the former communicates with the latter. One way to formulate this observation more rigorously is to say that lack of a context of mutual adjustment is common knowledge (Lewis 1969). A proposition *p* is common knowledge among a group of agents when each agent knows that *p*, and knows that each agent knows that *p*, and knows that each agent knows that each agent knows that *p*, and so on and so forth. Now let’s imagine two individuals, J and C, who enter into communication. J is a journalist and C is a news consumer. C knows two things: a) that they cannot consult J; b) that they believe (1). We may say that communication between J and C occurs in a context of mutual adjustment whenever (b) is common knowledge between J and C.

If it is common knowledge between journalists and news consumers that their communication occurs in a context of mutual adjustment, this fact affects also their mutual expectations. For instance, if the news consumer knows that the journalist knows that (b), the former may expect the latter to be more precise and unambiguous in their writing. On the part of the journalist this means, for example, that they are expected to make only their intended implicatures as accessible as possible, for there might be different but similarly plausible interpretations for what exactly a journalist conveys when making a statement. At the same time, this means that journalists must strive to avoid conveying, even unintentionally, unintended implicatures.

Notice that the same thing, to some extent, is taken into consideration in ordinary conversations too, for implicature indeterminacy is such a common element in every

communicative interaction. But within journalistic communication it carries much more weight. The reason, again, is that journalistic communication does not occur in a context of mutual adjustment. Since the news consumer cannot ask for clarification, the journalist has a far more pressing duty to avoid implicature indeterminacy.

Another difference between Gricean maxims and norms of news publication is that some attitudes towards the former may not be equally permissible when it comes to the latter. Overt violations of norms of news publication – what Grice would call "flouting a maxim" – are not always tolerated, especially with norms that resemble the maxim of Quality[54]. Think about jokes, teasing remarks, hyperboles, metaphors, euphemisms and many other cases of non-literal uses of language. As per Grice (although not for more up-to-date pragmatic theories) those cases are blatant/overt violations of the Maxim of Quality. Normally, speakers may resort to overt violations of the maxim of Quality for some reason. For example, I may deliberately say something false to make fun of my addressee.

But the same kind of violations may be forbidden in the context of journalistic communication, or at least not permitted with the same degree of freedom. Most of the time the use of irony to report facts, for instance, is unusual for a journalist, and it usually strikes the audience as alienating. This does not mean that non-literal uses of language are totally banned from a newspaper or a newscast. Opinion pieces can be rife with metaphors and hyperboles, for instance. But whenever we deal with reports of facts and events – the core of the information that we get from a journalistic source – it is convention to avoid the abuse of figurative language as much as possible, and to resort to it only when its use may help to convey more and better information than its alternatives[55]. For instance, it might be perfectly appropriate for a journalist to report that two Heads of State want "to sit at the table" to convey the idea that they intend to initiate negotiations. In this case, the metaphor effectively serves the purpose of delivering information, albeit in a more evocative way.

[54] The maxim of Quality is one of the four maxims that govern communicative exchanges between speakers, according to Grice (1975). These are expectations that it is rational to entertain with respect to speakers' contributions to the conversation. The maxim of Quality consists of a super-maximum – "Try to make your contribution one that is true" – and two sub-maxims: 1) "Do not say what you believe is false", 2) "Do not say that for which you lack adequate evidence".

[55] Genres also have an impact on the degree of admissibility of figurative language. Tabloid newspapers may be legitimised to make much wider use of it than, for example, newspapers that only cover news from around the world.

A third reason of discrepancy between norms of publication and Gricean maxims concerns differences in how expectations towards the former work. In particular, expectations towards norms of publication are mediated by the trust relationship we have with the information source. If we trust the source of information as a journalistic source, we will naturally pay more attention to the norms governing the formal aspect of news, giving them more priority. An interesting case of this mediation occurs for what in Grice's terms would be covert violations of the maxim of Quality: lies. How do we behave when a journalist deliberately reports falsehoods?

On the one hand, it seems that we, as readers, are more prone to assume that a norm analogous to the maxim of Quality is always satisfied in the context of journalism. The reason is that whenever I trust the fact that x is a journalistic source, I will presume that information on x is provided truthfully. In fact, when we refer to journalistic sources, like newspapers, we naturally take for granted that the author did not communicate something they believe to be false, given that journalism should always aim at the mark of truth. As a consequence, the observance of other norms (maxims) becomes a matter of *more urgency*: we especially expect journalist not to presuppose information unknown to the readers or that cannot be easily accommodated on the fly; to be perspicuous, unambiguous and brief; not to be be redundant, maintaining a balance between headlines and articles in terms of amount of information, and so on and so forth. It is not that these norms are in themselves more important than one concerning truthfulness. On the contrary, as we shall see in the next section, truthfulness stands on the top of the norms' hierarchy. Rather, it is more accurate to say that we pay more attention to the observance of these other norms, once we assume that x is a journalistic source.

This also helps to explain why, at the same time, we seem to have the opposite intuition about respecting a norm of truthfulness in the context of journalism. In fact, it seems that news consumers find lies of journalistic source more reproachable than when the same infringement occurs in talk exchanges: observance of a norm analogous to the maxim of Quality seems even more pressing for journalists. That is why we sanction more harshly journalistic mistakes, even accidental ones, and demand maximum professionalism from reporters.

I think that the reason behind this apparent irreconcilability of expectations – that is, expecting that observance of the maxim of Quality is at the same time less urgent and more urgent – is due to a difference in sets of journalistic norms which are taken into consideration each time. Expecting that (something analogous to) the maxim of Quality is being strictly observed is probably something that pertains more to norms of news production. The scrupulous verification of information, already mentioned by Pepp et al. (2019), is a good example of that. That is, we somehow believe that a piece of news whose informative contribution was poorly verified should never appear in the column of the newspaper to begin with. But as long as it appears in that column, we take for granted that norms of news productions have been observed and that the author of the article is truthful. This, again, is because we start from the assumption that that column is part of a journalistic source that we trust. Then expectations towards norms of news publication kick in and, suddenly, observance of other norms becomes more pressing. But this does not imply that observance of truthfulness was irrelevant: it was simply bound to a previous stage of journalistic activity[56].

Finally, another reason of diversity with Gricean maxims is due to the importance of style in journalistic writing. With "style" I do not only mean the syntax, the grammatical elements and the overall sentence complexity of journalistic production. Rather, the label also includes the whole set of tacit rules and conventions that dictates how journalists should craft their works.

For instance, a famous rule of thumb for journalistic writing is the so-called "inverted pyramid scheme" for articles. This rule simply dictates to order information that the journalist want to convey through an article according to its importance: substantial information should be at the top of the article, while secondary one should follow. Since the primary goal of a journalist is to deliver accurate and newsworthy information, the inverted pyramid scheme allows the news consumer to interrupt reading at any point without significant losses of information.

Another dissimilarity in style concerns objectivity. Objectivity may mean two separate things, both equally important for journalists: objectivity as impartiality and

[56] Other norms of news production may work in the same way. Consider the commitment to the interest of readers. This norm entails, among other things, a duty of selecting only news that might interest to the reader (that is, newsworthy information). So when an event is not reported by news media outlets we automatically judge that it was none of our concern, and vice versa.

objectivity as personal detachment. The former means framing facts in a neutral way, with the intent of not conveying sneaky value judgments. For instance, a non-objective journalist in this sense that recounts the presidency of Ronald Regan may write (1):

(1) Ronald Reagan's leadership style was inspiring and visionary, setting a standard that no other president has matched.

While the objective version of (1) may be something like this:

(1*) Ronald Reagan's leadership style is often noted for its emphasis on communication and a strong national defence, and he implemented policies that aimed to reduce government regulation and taxes.

In (1*) the subjective descriptors, "inspiring" and "visionary", are removed. Instead, the focus is shifted to factual aspects of his leadership style, such as "emphasis on communication" and "strong national defence", which are more neutral and can be substantiated with evidence. Besides, phrases that express personal judgment, like "setting a standard that no other president has matched", are taken off too. The objective version focuses on specific actions and policies president Regan is known for, like "reducing government regulation and taxes", which can be objectively verified without implying that he set an unmatched standard.

The latter meaning of objectivity entails the attitude of someone who refrains from express their point of view when reporting facts. To do that, journalists usually write in the third person (he, she, they), instead of using first-person pronouns (I, me, my, our, we), to maintain distance from the covered facts. They focus on what they observe and hear, net of personal biases, rather than on what they think, believe, or feel subjectively.

Gricean maxims do not ask for objectivity in no way whatsoever. The reason is that, for the smooth performance of an ordinary conversation, is not necessary that the speaker assumes an objective point of view in reporting facts, be it framed as impartiality or personal detachment. Instead, this further requirement is quite pressing for journalists. The stylistic features mentioned so far, serve as hallmarks for the social category of journalists and establish another ground for our expectations towards a journalistic product. Hence, we do not just expect journalists to observe norms that ensure an effective informative exchange, but we also expect journalistic headlines and

articles to have certain distinctive characteristics. The set of journalistic norms of news publication is therefore broader and more sophisticated than the one of Gricean maxims.

For all these reasons, even though some of the principles related to truth, clarity, modulation of information, etc., bear a strong resemblance to Gricean maxims, the norms of journalism neither correspond to nor can be reduced to Gricean maxims. Moreover, the former, unlike the latter, do not lend themselves to strategic violations. Grice, in fact, hypothesized that a certain portion of what we communicate is conveyed implicitly, generated by the apparent violation of one or more maxims, but which, under the assumption of the cooperative principle, turns out to be an exploitation rather than a violation of the maxims. This, however, does not occur in journalism. For a journalist, it makes no sense to violate the norm that requires them to be objective in order to strategically convey a particular type of content; this simply cannot happen. This is partly due to the fact that, in the absence of a context of mutual adjustment, there is no genuine dialogue between the journalist and the news consumer that would allow for the type of maxim exploitation found in Grice's dialogic examples, where one speaker asks a question, and the other responds by openly violating a maxim.

2.2 Hierarchy of norms

Not all norms of journalism are of equal importance. Some of them have more, or less, weight than others. Consequently, in many cases the violation of a specific norm may be more severe than the violation of another one. What I am suggesting is that there exists a sort of hierarchy among norms of journalism, showing that some of them have always 'priority' over others. The existence of a hierarchy helps us explain that violations of journalistic norms do not always prefigurate a case of fake news, for the violation, according to FN_A, must be significant. Let's see why.

As we saw, due to its institutional nature, news has the deontic power to instil a set of expectations in the news consumer. Whenever a news item falls short of those expectations, the news consumer may feel entitled to reproach the journalist for not being compliant with their obligations (which are inscribed in the norms of journalism). But those obligations may vary in importance. For instance, we may reasonably say that a journalist who reports something they did not verify is more severe than another journalist who is sloppy in the application of the inverted pyramid scheme. We have the

strong intuition that the observance of the norm of news production that requires scrupulous verification of information is *more urgent* than observance of the norm of news publications that requires adherence to the inverted pyramid scheme. As a consequence, news consumer may feel more entitled to reprimand a journalist that violates the former norm than another who violates the latter norm, ceteris paribus. This points to the fact that a hierarchy of norms can possibly exist[57].

In general I maintain that, ceteris paribus, norms of news production have priority over norms of news publication. Sticking to the last example can help us to accept more easily this conclusion. In fact, I have the intuition that a source of information who repeatedly publishes unverified reports goes more against the profession and the values of journalism than another source that verifies information but gets it across sloppier than the former one. In fact, sloppy communication, while problematic, does not inherently violate the responsibility of a journalist in the same way as failing to verify facts. A journalist who reports unverified information violates the epistemic responsibility to "get things right", a more critical obligation than presenting information elegantly: sloppiness in presentation can obscure meaning but does not inherently introduce falsehoods into the public discourse. Since norms of news production aim at the preservation of epistemic goods, while norms of news publication prescribe behaviours for the optimal transmission of those goods, it seems that observance of former norms are more urgent than observance of the latter norms. Therefore, at least in principle, violating norms of news production is usually more severe than violating norms of news publication (as I clarify below, things are not always so simple).

Besides, norms of journalism can be ordered according to their importance even within their own category. For instance, it seems legitimate to think that adherence to the norm of news production that requires a piece of information to be potentially interesting for the public (newsworthiness) is more urgent than adherence to the norm of news production that requires proportionality in the amount of coverage. The same goes for norms of news publication. Making intended implicatures as accessible as possible

[57] Also the list of obligation offered by Kovach and Rosenstiel (2014) is presented following an order of importance. For instance, the obligation towards the truth should come first for journalists; everything else, however important, follows this first principle. This proves that cues for a hierarchy of norms are present within reflections of distinguished journalists too.

for the news consumer – and avoiding conveyance of unintended ones – seems more urgent than making the text brief. Sticking to these two examples, violating the former norm is more severe than violating the latter norm.

Assessments of the severity of violations also gives us the opportunity to explain an aspect of FN_A and FN_B that has been neglected so far: the degree of violation. In short, we do not want to consider fake news also cases in which the violation of a norm of journalism is partial or negligible. For example, we do not want to consider as fake news a piece of news that is somewhat obscure in one sentence but otherwise perspicuous. It would be an unfair judgement. Both FN_A and FN_B keep track of this necessity. FN_A states that news is fake when it does not live up to our expectations *to a sufficient degree*, while FN_B states that news is fake when it *significantly* violates norms of journalism. Sure, the metric is once again subjective: we cannot pinpoint a strict, mathematically determined threshold beyond which we are sure that a full-fledged violation (or disappointment of expectations) has occurred. But comparative assessments can be made, and they are sufficient for our purposes.

For example, violations of norms higher up the hierarchy are more significant than violations of norms further down the hierarchy. Almost all the cases analysed so far in this paragraph are just that. They measure the degree of violation based on the importance of the violated norm(s). It is a qualitative measure of violation: the more important a norm, the more significant it is to violate it. But this is not the only yardstick at our disposal. There is also a quantitative measure. In fact, especially for norms of news publication, there are violations of the same norm that are more pronounced than others. To recall an example made above, there is difference between an article which has only one sentence that is obscure or ambiguous and another article which is totally unintelligible. The norm at stake may be the same, but the degree of violation is quantitatively different.

Following the same principle, we may say that deeper and more pronounced violations of norms further down the hierarchy can sometimes trump negligible violations of norms higher up the hierarchy. For instance, the issuance of a completely unintelligible news item may sometimes be more severe than issuing news items without disclosing conflicts of interest. Normally, we may say that to publish news that is biased or lacks transparency about conflicts of interest of its producers is more severe

than publish news that contains some obscure sentences. That is, the norm of production that requires transparency about personal, financial, or institutional interests that could influence the way a story is recounted is usually more important than the norm of news publication that requires perspicuity of sentences.

Imagine a news channel, owned by a big oil corporation, that consistently downplays the effects of climate change or questions the scientific consensus about global warming, without disclosing their corporate ties. Apart from ethical considerations, the lack of transparency about conflicts of interest may undermine the trustworthiness of the news itself. Readers might believe that the information about climate science issued by that news channel is impartial when, in fact, it is biased due to hidden motives. On the other hand, a story with a bit complex jargon or some convoluted sentences can be hard to follow, but if the facts presented are accurate, the problem is more about clarity. While less intelligible sentences are problematic in the short term, because they reduce accessibility, their long-term consequences are generally less severe. Poorly written content might lead to reader frustration, but it typically does not have the same long-lasting impact on the reputation or the integrity of journalism.

Nevertheless, sometimes very pronounced violations of the norms of news publication may weigh more than light violations of the former. Consider an article about a public health crisis, such as an outbreak of a contagious disease. The article's purpose is to inform readers about the situation and provide critical guidance on how to protect themselves and others. A biased version of this article might downplay the severity of the outbreak or emphasize a particular political perspective on the government's handling of the crisis. While the bias could influence how readers interpret the situation – for instance, leading them to minimize the risk or criticize certain public health measures – it would still likely communicate key facts, such as the symptoms of the disease, how it spreads, and what actions to take to mitigate risk. News consumers, even if they sense the bias, can extract the necessary information to make informed decisions.

Now imagine the same article, but instead of being biased, it is poorly written and incoherent. The sentences are disorganized, the terminology is inconsistent, and the structure of the article makes it difficult to follow. For example, the symptoms of the disease might be mentioned in one paragraph, contradicted in another, and buried under

unrelated details. The preventative measures might be described in technical jargon without clear explanations, leaving readers unsure about what steps to take. In this scenario, even though the article might contain unbiased information, the lack of clarity and coherence renders it almost useless. Readers may be unable to discern how to protect themselves or others, leading to confusion, panic, or inaction.

Extremely poor writing can block access to the information in the first place, making it impossible for the reader to engage with the content or use it for informed decision-making. In this sense, poor writing can be considered a more severe infringement on journalistic standards in situations where the stakes are high, such as in public health reporting, because it compromises one capability of journalists – to communicate critical, clear, and understandable information to the public. Therefore, sometimes violations of norms of news publication may be more severe than violations of norms of news productions.

To recap, we have hierarchies among categories of norms and hierarchies within a single category of norms. First, observance of norms of news production matters more than observance of norms of news publication, ceteris paribus. This means that, usually, violations of the former category of norms are more severe than violations of the latter category. Furthermore, norms within the same category can be ordered in a hierarchy according to their importance too; so that we can compare the severity of violations of two norms belonging to the same category (production or publication). Finally, the degree of violation is also important: if there is a very pronounced violation of a lower-ranking norm, this may outweigh in importance a slight violation of a higher-ranking norm in specific contexts.

This explains why, even though one might be tempted to draw a sharp line between low-quality journalistic products and fake news, such a clear boundary can be elusive. In reality, the norms of journalism – be them about the production or publication phase – may be violated to varying degrees, creating a continuum rather than a simple yes/no condition. A news outlet might cut corners only sporadically or exaggerate headlines in a way that remains barely consistent with minimal journalistic standards. Conversely, another might continuously and substantially fail these norms, pushing its content decisively into fake news territory.

On this view, then, the boundaries between low-quality journalism and fake news are mobile; not every violation immediately topples a piece of journalism into outright "fakeness", nor does every sloppy piece of reporting become fake news. Rather, the concept of fake news picks out cases where the defects in upholding journalistic norms are significant enough to undermine or forfeit the news status altogether. That threshold can be flexible or context-dependent, reflecting how audiences judge the severity and frequency of norm breaches.

3. The institutional account of fake news

In the first chapter I said that current accounts on fake news may be roughly divided in two groups: the conservatives and the privativists. Conservatives assume that fake news is still a kind of news, while privativists deny this. The idea of conservatives is that news is a kind of information – viz. recent reported information. Whereas the reason that leads privativists to the opposite conclusion is that news is more of a status that a piece of information may or may not deserve. Fake news does not have the right history of production so it is not really news, even though it can look very much like real news. So what is my position on the matter? Is fake news still news or not in my account?

At first glance, my account seems to be more compatible with the privative response. In fact, I said that news indicates a status function, and specifically the one that a piece of information gets whenever it is issued by a journalistic source. So it seems natural to think that information either has this status or it does not. If it does not have this status, then we are dealing with fake news. But the question is not so simple. In fact, framing the possession of a status as an all or nothing issue comes from a simplistic way of conceiving the assignment of status functions.

A first problem of this frame is that there are many ways through which a piece of information may not get the status of news. Consider the example of tree rings. Tree rings are an invaluable source of information for scientists. They do not give us only information about the age of trees but, among other things, they also offer indirect evidence about the temperature levels in the past. Notwithstanding its value, information that we get from tree rings is not news. Sure, the reason is blatant: tree rings are not a journalistic source. But as far as I know, no one has ever been fooled into thinking that tree rings deliver news. In contrast, every day so many people mistake

online information issued by non-journalistic sources for news. If we want to keep our discourse to phenomena that really interest us, like (fake) news that we encounter online, we should separate them from other, irrelevant cases in which information does not get the status of news.

A simple way to bypass this problem is to reflect on the concept of news. The status function of news is regulated by a structure of constitutive rules that goes down to the ones that regulate the profession of the journalist. Constitutive rules, as we now know, are rules of the form "X counts as Y in the context C". By bestowing the status Y to the object or the person X, we are conferring deontic powers to the latter. Deontic powers are either enablements or requirements. For example, a particular position (X) in chess (C) is considered as checkmate (Y). The status of checkmate *enables* the player who reached that position to win the game. At the same time, that constitutive rule *requires* that in order to win the game a player must reach the position of checkmate, and not any other. Whenever we recognise that this rule, plus many others, is in force and we act according to its enablements and requirements, then we are playing chess. When we accept that those rules are in force, and we also know their enablements and requirements, but we nevertheless decide to not comply with them, then we are violating the rules. If I checkmate my opponent but I decide not to stop and take the victory, I am violating the rules of chess.

But what happens when one does not act according to the deontic powers because they cannot recognise that the rules are in force? Imagine, for instance, that two kittens left alone in the house are moving chess pieces haphazardly on the chessboard as part of their play session. Are they violating the rules of chess? Probably not, the reason being the kittens are not even in the business of observing chess rules. They do not recognize the existence of deontic powers, because they cannot perform the assignment of a status function to begin with. There is no collective agreement to assign status functions to inanimate objects like chess pieces.

Therefore, I find it useful to distinguish between violating a rule and failing to recognize it in the first place. In considering fake news, we should pay attention only to the first type of occurrences. We are not concerned with cases of information that do not deserve the status of news for reasons that have nothing to do with the respect of the deontic powers assigned by the proper constitutive rule. Information that we get from

tree rings belongs to the second type of occurrences. It does not deserve neither the status of news, because it does not fulfill the requirements of a journalistic source, nor of fake news, because we would not expect this to be the case. We do not expect the kittens to comply to the rules of chess when they are rattling chess pieces, so blaming them for violating those rules does not make sense. Blaming tree rings for not observing norms of journalism would be equally absurd.

The case of information that we find on newspapers, instead, is different. When we are reading a trusted newspaper, we assume that information that we get from it is news. In other words, we have the expectation that what we are reading is news. If those expectations are severely disappointed, then we get fake news. But an initial expectation must be formed to begin with; otherwise there is no disappointment. I cannot be disappointed by the fact that the kitten does not move the knight in a 'L' pattern when it slams the knight with its paws on the chessboard. In contrast, I get upset and annoyed if the author of a journalistic article did not verify what he reported on the newspaper.

Only violations of norms of journalism *in cases where we expect them to be observed* identify fake news: any other case that entails lack of the status function of news is not captured, because in order to violate a norm you have to first recognize its existence and its normative power. Information that we get from natural sources does not violate norms of journalism, because there is no acknowledgement of journalistic norms in the first place. And this observation leads us to the central question of this section. Since fake news may arise only from cases in which we expect the observance of the norms of journalism, I maintain that fake news is information that possesses the status of news, albeit undeservedly. Fake news is news to which we have undeservedly conferred that status. And the reason we did this is because we expected the source of that information to be a journalistic source, namely one that fulfils the obligations of journalism. Let's delve deeper into this aspect.

The constitutive rule for news establishes the condition for the conferment of the status of news. Whenever a piece of information is issued by a journalistic source, that piece of information is news. This entails that when we consume a piece of information that comes from (what we believe is) a journalistic source, then we are prone to conceive it as news. So there is always an initial stance in which the piece of information – regardless of any violation of norms – is news. The reason is that the

conferment of the status according to the rule is a matter of collective agreement, and not of personal decision. This is another reason why thinking of the conferment of status as an all-or-nothing thing is a simplification. In reality, often times there is no actual moment at which the status function is bestowed, as in a kind of initial baptism. Members of a community are not always directly involved in the conferment of status; collective agreement does not just work in this way. We must not make the mistake of reducing status assignments to the most obvious cases, such as weddings or graduations. Those are just evident manifestations of a phenomenon that usually goes unnoticed. There is a bunch of reasons for this.

For instance, most of the time we are incapable of tracing the origin of an institution and its constitutive rules. Take again the example of money. We now know that using a universal medium of exchange to facilitate trades among humans was a smart move. Through the collective assignment of a handful of status functions we ended up solving practical issues of massive importance: the impossibility of both splitting many commodities (think about the cattle in ancient times) and setting a general reference value for our goods. But how did we accomplish it? And when did this turn precisely occur? Historically, there is really no exact point at which human communities came together and decided to assign the status of medium of exchange and store of value to pieces of metal. This happened slowly, over time, in a process full of internal upheavals: barter gave rise to commodity money (that is money with intrinsic value, like gold coins), commodity money then led to convertible paper currency, and the currency subsequently gave way to contemporary fiat money (Dowd 2001). But the whole process was for the most part unintentional, it was not really guided by collective decisions such as a show of hands vote. The assignment of the status of money to metal coins, at least in the beginning, took off as a convention without awareness.

Another fact is that we usually accept the existence and the maintenance of many institutions simply because we take them for granted. For example, we take it for granted that metal coins have value as money because we were born into a world where the institution of money has long been in place. We have never really been asked about this. I did not express my agreement (or disagreement) in bestowing the status of money to, say, the metal coins distributed by the Italian mint simply because none questioned me about it. Yet, I accept euros as medium of exchange in my economic transactions.

Again, there was not a moment in the past where I decided whether to accept or decline the conferment of the status of money to metal coins coming from the Italian mint. I just passively accepted that.

Therefore, returning to the issue, initially information issued by a (supposed) journalistic source is basically always news, just as a banknote issued by the state mint is always money. Then whether that banknote has production defects is an individual assessment of that individual token that always comes temporally after the status has been assigned. And the same goes for news. Fake news, in my view, is still news because it was issued by a source of information that we consider journalistic; and yet that piece of information did not deserve that status, because it significantly violates norms of journalism and, which is the same, does not live up to a sufficient degree to realistic expectations towards journalistic product. So it becomes clear that *the real issue in fake news is located in the source of information that issued it, and not in its representational content*. Fake news is not a matter of false or misleading information, it's about deviations from the work standards of journalists, who have rights and obligations that descend from the status function of their own profession. Fake news is news that comes from a journalistic source that is not fulfilling its duties, either because it was not a real journalistic source to begin with or it is a real one which on that specific occasion failed in its obligations. Most of the time, fake news comes from the former case: sources of information mistakenly treated as journalistic ones.

Therefore, in a sense my proposal on fake news wants to be a compromise between the conservative view and the privative view. On the one hand, my account is conservative for eventually fake news is still news in my analysis. On the other hand though, contrary to existing conservative accounts, my definition of fake news is based on the notion of status function, which is more familiar to the privative faction. With "fake news" I do not refer to pathologies of information, nor to examples of insincere communication; instead, I refer to articles, news reports, documentaries and any other journalistic product that actually does not deserve that status. Fake news is news issued by sources which act like representatives of the institution of journalism but that in reality are not fulfilling the obligations that come with membership in that institution[58].

[58] Recall that institutions are systems of constitutive rules, while "obligations" here refer to the requirements (one type of deontic powers) that descends from the adoption of one or more constitutive

So contrary to the privative view, it is not the piece of news that is really passing itself off as news, but it is the source of information which is masquerading as a journalistic source. We might say, in a sense, that is the source who is really counterfeit, and not the piece of news. That is why my account is not committed to the same interpretation of the adjective "fake" in "fake news" that we find in the privative view of fake news (below, §3.1).

So far, we have answered the four questions raised in §1. We now know the kind of realistic expectations at stake in the definition of fake news, and how a piece of news may not live up to them; the norms of journalism, their typology and their internal hierarchy; how fake news violates expectations on journalistic product in a significant way; why fake news is still news, in a sense. In the following two subsections, I clarify my interpretation of the adjective "fake" in "fake news" and the theoretical features of the new concept of fake news that I have just stipulated here.

3.1 The interpretation of "fake" in "fake news"

We already know that the privative view buttresses the intuition that fake news is not news by resorting to the grammar of the adjective "fake". According to their interpretation, the expression "fake news" should be compared to the one of "fake gun". Since fake guns are not really guns, when we dub something as fake news we are claiming that something is not really news. But there are grounds for saying that this conclusion is disputable (Grundmann 2020). First of all, we may dispute the necessity to trace the meaning of "fake news" back to its ordinary use to begin with: as philosophers, we are entitled to use "fake news" – or whatever notion we would like to investigate – in a technical sense, without minding the way the notion itself, or one of its components, is deployed in ordinary language (this is what I do with "news" in the second chapter, for instance). And even though we still accept the idea of taking into consideration the ordinary use of a notion, we will find that there is evidence that goes against the very idea that "fake" in "fake news" behaves as in the expression "fake gun". This has to do with the tricky nature of privative adjectives.

Among the many problems linguists encounter when classifying adjectives, one is the question of whether a modifier such as "false" always acts as a privative or not.

rules. So adherence to the system of constitutive rules called "journalism" comes with the acceptance of the list of requirements/obligations that stems from it.

Several observations fall short with the interpretation of “fake” as a purely privative adjective. Consider for instance the sentences below from Partee (2021):

2a) A fake gun is not a gun.

2b) Is that gun fake or real?

This example strikes us out of the fact that we cannot readily explain the evident tension between the apparent truth of (2a) and the undeniable well-formedness and interpretability of (2b). In other words, if fake guns are not guns why does (2b) sound legitimate to ask?

Besides, consider the adjective “fake” in combination with other names, like “smile”. Should we say that a fake smile is not a flat-out smile? In contrast, it seems that, at least in this case, a fake article is still a smile, thus causing “fake” to be a subsective adjective, since the inference pattern “if X is Adj(Np) then X is Np” holds for “fake smile” (recall that subsective adjectives indicate adjective whose denotation in combination with a noun is a subset of the denotation of the noun). And the same logic applies to other cases, as was the case with “fake video” and “fake website”. Fake websites are real websites with fraudulent content, whereas fake videos are real videos which have been doctored or manipulated in their content.

The same dual behaviour can be found in other supposed privative adjectives, like “false”. Take the example of authors who endorse a factive interpretation of information. As Paul Grice once famously wrote concerning his maxim of Quality, «false information is not an inferior kind of information; it just is not information» (Grice 1989: 371). In that context Grice was straightforwardly affirming the privative nature of “false” when the latter is coupled with the noun “information”. Along the same lines, Dretske affirmed that «false information, misinformation, and (grimace!) disinformation are not varieties of information – any more than a decoy duck is a kind of duck» (Dretske 1983: 57). But then consider the case of “false alarm”. It seems equally true to say that a false alarm does in fact belongs to the extension of the name “alarm”. After all, an alarm signalling a non-existent danger is still a full-fledged alarm.

What is the point of this detour on adjective classification? It seems that something very similar happens in the case of “fake news”. In the case of “fake smile” or “false alarm”, the expression “fake news” does not necessarily entail a privative behaviour on

the part of "fake". The reason is that news, contrary to guns and information, *is not an object that can be reduced to one essential function*. A gun is an object that might perform an array of different functions: I can intimidate others with it, I can use it as a blunt object and hit other people etc. But a gun that fails to shoot presumably lacks its essential function. A piece of information is a semantic entity, something that has the property of representing the world as being in a certain way. Regardless of whether you think that information incorporates veridicality, as proposed by Grice and Dretske, a piece of information which fails to represent a state of affair cannot really be information. Information without semantic content is not information according to both factive and non-factive accounts of information.

The opposite happens for nouns like "alarm" or "smile", whose denotata in the real world are devoid of one distinctive property or function. For this very reason, when "fake" is coupled with nouns like "gun" – that is, nouns that point to objects, with a recognizable primary function – the former usually performs a privative role. But everything changes when "fake" is attached to nouns along the lines of "smile" and "alarm"; and I maintain that this is precisely the case of "news". In "fake news", the adjective "fake" does not deny the membership of the nominal phrase to the general noun category, because news is not an ordinary object with a primary function. News is a name for a status function that we may assign to information, and we now know that status functions are entities that exist only through collective agreement. They are special precisely because their existence do not depend on the characteristics of the object to which they are bestowed[59]. So the status of news is not affected by semantic or epistemic properties of information. Unlike information as interpreted by Grice and Dretske, for instance, veridicality is not an essential property of news: false news is still news for all intents and purposes. So the adjective "fake" in "fake news" might not perform a privative role; and this conclusion is shared by other scholars as well.

Based on observations about adjective-noun pairs, Cappelle, Denis and Keller (2018) have hypothesized that the interpretation of "fake" is context-sensitive, and depends on the kind of noun to which "fake" is attached to. To test this hypothesis, they have extracted the 100 most frequent fake-nouns combinations from the English Wikipedia

[59] Recall that the constitutive rule for news is not a definition of news (above, p. 89), since constitutive rules are just formulas that remind us how we ordinarily assign the status in question. They make explicit what we normally do tacitly and unconsciously in everyday life.

2013 corpus that contains a total of 1920 occurrences of fake-noun, among which there is also the combination "fake news"[60]. Then, they have classified these 100 combinations according to seven different interpretation of the contribution of "fake". Of these seven, three are privative interpretations, other three are non-privative and the last one is ambivalent (for some interpreters is privative and for others is not). After that, the authors assigned manually the 100 combinations to each of the seven interpretations. According to their assignation, the combination "fake news" appears under the interpretation no. 5, which is one of the non-privative interpretations. Together with "fake article", "fake blog", and others, Cappelle and colleagues claim that "fake" in "fake news" conveys the information that fake news is still news plus one or more of these additional non-privative interpretations: 1) it did not come into being normally; 2) it was made to look like normally originated news; 2) it appears true but it is not; 3) it has one or more of the effects of normally originated news.

Afterwards, the team tested whether the different interpretations, including the no. 5, assigned manually to fake-noun combinations were not random but correlated with semantic properties of the nouns involved. In other words, they wanted to understand if nouns with similar semantic properties yielded similarly interpreted fake-noun combinations. To test this, the study used a distributional semantics approach, which analyses the contexts in which words appear to understand their meanings. Specifically, the study looks at word embeddings – mathematical representations of words based on their usage in a large corpus (Wikipedia in this case). To understand how similar the nouns are to each other, the researchers used a Gaussian kernel, a mathematical function used to model local deformation in computer science, to transform the distances between them. This method made the differences between the nouns more noticeable. If two nouns were very similar, they got a high score; if they were very different, they got a low score. Cappelle and colleagues decided that two nouns would be considered connected if their similarity score was above a certain level (0.01). They chose this level carefully to make sure the network of nouns was well-connected, but not too cluttered.

[60] It is important to note that the Wikipedia corpus considered here precedes the 2016 US election and subsequent diffusion of the expression "fake news" among journalists and social media users. That is why Cappelle and colleagues have not found as many occurrences of this pair as expected in the Wikipedia corpus. The study also does take into consideration a possible twist in the meaning of fake news after the 2016 US election (e.g. the weaponization of fake news).

Next, they wanted to create a visual map of how similar the nouns were to each other. They used an algorithm called "ForceAtlas2" (Jacomy et al. 2014) to do this. This tool works like a system where similar nouns pull towards each other like magnets, while less similar nouns push away from each other. Finally, they checked how well the clusters of nouns in this visual map matched the different meanings of "fake" that they had identified earlier. They used a method called "performance P" (Fortunato 2010) that helps determine whether nouns that have been manually assigned the same type of fake-noun interpretation are actually grouped together (i.e., connected by an edge) in the distributional network, and whether nouns with different interpretations are not grouped together. Basically, a higher performance P would indicate a better correspondence between the network and the manual classifications.

The result was in line with the starting hypothesis: Cappelle and his team achieved a performance P of 0.78, which is significantly better than random chance. This means that nouns with the same fake-noun interpretation were generally linked by an edge, while those with different fake-noun interpretations are usually not linked. This result indicates that the interpretation of a fake-noun combination is influenced to a considerable degree by the semantic properties of the noun, thus supporting the initial hypothesis on context-sensitivity of the adjective "fake". To sum it up: we cannot take for granted that "fake" in "fake news" has a privative behaviour, as many authors of the privative view accept. It can perform a different role.

The main idea of Cappelle and his team is that the general semantic contribution of "fake", both in privative and non-privative interpretations, is to indicate that the object denoted by the noun is somehow *deceptive*. In cases such as "fake news", one of the possible interpretation is that the object in question – here the piece of news – might not have come into being like ordinary objects of that type, thus deceiving about its process of manufacturing[61]. And this interpretation is fully consistent with my proposal. In fact, fake news is news that significantly violates norms of journalism, where journalistic

[61] Cappelle and colleagues arrive to this interpretation drawing on the notion of qualia structure, introduced by Pustejovsky (1995). Qualia structures are semantic dimensions of nouns and nominal phrases that represent properties of the object denoted by the noun/nominal phrase. According to Pustejovsky, there are at least four qualia structures: 1) Constitutive, that indicates the relation between the entity and its parts, what it is part of, or what it is made of; 2) Formal, that which we use to distinguish or pick out the entity (like its shape); 3) Telic, the purpose of the function of the entity; 4) Agentive, factors connected with the origin or way in which the entity came into being (this summary is taken from Del Pinal 2015). In interpreting combinations like "fake news", Cappelle and colleagues believe that "fake" affects the *agentive qualis* of "fake news".

norms represent a list of both obligations and distinctive epistemic and communicative practices of the profession of journalism, the observance of which confers the status of news. Fake news is news produced or published in severe violation of norms of production or publication (or both), and so it is news that did not come into being as it should have been. As a consequence, if they cannot find evidence of a violation, fake news can deceive the news consumer into thinking that source of that information is acting as a journalistic source, when actually it is not (either because it was not a journalistic source to begin with or it performed poorly in this occasion).

This last observation is crucial, for it is easy to fall again into the temptation to think of fake news deception as being based on imitating the appearance of real news. This is not the kind of deception that I want to stress here. Concerning the idea of deception conveyed by "fake" across privative and non-privative interpretations, Cappelle and colleagues write that this consists in the fact that something/someone is deliberately presented as something/someone that looks like something/someone else in order to mislead. But in fake news what is presented as something that looks like something else is not the piece of information, as is the case for privativists. *It is the source of information*, the object that fills the C-place in the constitutive rule of news, that is passing itself off as compliant with the institution of journalism, when actually it is not. Fake news is an attempt to mislead the news consumer about the identity or the reliability of the source that issued that piece of information. It is a way to feign the legitimacy of real journalistic sources in a world where an official codification of the institutional entity of news does not exist. This is the core of the institutional account of fake news.

3.2 Properties of fake news

The concept of fake news that I have been describing so far presents interesting features that needs to be highlighted. Some of them are already captured by other definitions of fake news on the market, especially from the privative view. Still no characterization of fake news presents all these features together.

The first prominent feature of my definition of fake news is that it is conceptually separated from pathologies of information, like misinformation or disinformation. The reason is that we have disentangled the concept of news from the one of information, so

that the former is no more a species of the latter. News is not information that has been reported by someone to a minimum number of people, but it is a status function that may supervene on information. Accordingly, fake news is not really a flawed kind of information but points instead to a news item that does not live up to our expectations towards journalistic products to a sufficient degree (or which significantly violates norms of journalism). Therefore, fake news may well not have the epistemic flaws that misinformation or disinformation has.

For instance, fake news can safely be true in its content. That is, if the piece of information upon which the status function of news has been undeservedly assigned is true, we can still have fake news. In fact, the journalist in question may have reported the truth by accident, and in doing their investigation they may not have followed the norms of journalism. For instance, they could have published a news item that was pulled out of the air. Jaster and Lanius offer an interesting example of this kind:

> Suppose *Russia Today* fabricates a news story about Hillary Clinton: according to the story, Clinton committed tax fraud. The authors of the report take it to be false and propagate the story with the intention to deceive and spread disinformation. As it happens, however, Clinton has in fact committed tax fraud. The story turns out true; it is neither false nor misleading. Is the *Russia Today* report fake news or not? (Jaster and Lanius 2018: 218)

Jaster and Lanius answer negatively to their question: *Russia Today* has not published a piece of fake news, but it merely attempted to do so because the piece of news was true in the end. Croce and Piazza (2021) have brought compelling arguments against this conclusion. The main one highlights the analogy between this case of *Russia Today* and the classic ones invented by Gettier (1963). In Gettier's scenarios, there are epistemic agents who have a justified true belief in a proposition, yet fail to have knowledge due to an accidental relationship between the available evidence and the truth of the relevant belief. The example of *Russia Today* seems to have the same structure: news consumers who watch the Russian state-controlled news television network would acquire a true belief by consuming a story that was not grounded on evidence; thus making the connection between the belief of news consumers and the truth too shallow to be knowledge. But then, Croce and Piazza affirms, this is exactly what we should expect

from fake news: that by consuming it one should not be able to acquire knowledge. So the piece of news spread by *Russia Today* is fake news for all intents and purposes.

I believe that my account of fake news fits well with this critique of Croce and Piazza. Recall in fact that observance of norms of journalism ensures the epistemic value of the journalistic product: readers expect that from consuming a news item they are able to acquire knowledge. Therefore, news items that are accidentally true – because they do not comply with norms of journalism that ensure quality of the outcome – do not deserve the status of news. Again, it is a matter of expectations and not certainty: at the end of the story, the genuine news item may well be false.

This observation leads us to a further point: in my account, false news is not necessarily fake news, and vice versa. The reason is that violation of the norms of journalism does not necessarily entail lack of veridicality. In fact, even though one of the primary norms of news production says that the first obligation of journalism is to the truth (above, p. 92), this does not imply in any way that the news content must necessarily be true. What is required to a journalist is to always aim for the truth, not necessarily to reach it every time. So failures in reporting the truth are a common phenomenon in journalism, and they must be expected. However, if the journalist has scrupulously followed the norm of their profession and they stumbled into error despite everything, then we should not dub their product as fake news. We do not necessarily devalue the work of a journalist that accidentally reported a falsehood: bona fide errors do exist, and sometimes misfortunes get along the way of good journalism too[62]. That is why there is usually a certain tolerance threshold for these errors. For example, we are willing to accept a false report when this is followed by prompt rectification and/or an apology from the newspaper or the journalist who made the mistake, usually complete with an explanation of the origin of the error.

These acts of 'reparation' are important because their mere existence points to the normative force of the norms of journalism and the obligations for which they stand for. Whenever a journalist publishes a fake news item, they are falling short of the expectations that readers have of their products. It is like breaking an implicit pact of

[62] This argument mirrors the critique moved to Jaster and Lanius (2018), when they support the addition of an intention to deceive into the definition of fake news in order to exclude cases of accidental mistakes (above, p. 29). It is not a matter of intentions; what is at stake here are the expectations that we have towards the obligations of a specific profession, the journalist.

trust that exists between news producers and news consumers. This breach of trust is another way to describe the failure of the journalist to fulfil the obligations of their profession. And as we shall know by now, this gives to the reader the right to reproach and criticize the journalist for that. By apologizing and rectifying the error, the journalist (or the newspaper) is simply trying to reconstitute that pact. Here's a real example of that. In an article by Sullivan (2013) named "Repairing the Credibility Cracks", *The New York Times* recalls the troublesome story of Jayson Blair, a young Times reporter that lied and fabricated information multiple times for years to the detriment of the newspaper. After some time, when the editorial staff began to be suspicious towards Blair, even thanks to audience reports, the newspaper started an internal investigation that in 2003 led to the dismissal of Blair as well as *The New York Times*'s executive editor and managing editor. The newspaper itself was the first to reveal the story, even at the cost of losing part of its audience, and make amends, all with the intent of «repairing the credibility cracks».

Instead, by explaining the origin of the error, the journalist is giving us evidence of the fact that they did observe the norms of journalism (or at least they tried), but nevertheless failed to deliver accurate information. This evidence is important for the news consumer, because they can adjust their evaluation in response to an initial disappointment of their expectations. If they initially deemed the report as fake news, they now might change their mind due to new evidence that testify how the journalist actually tried to fulfill their obligations, notwithstanding the poor result.

But the breach of trust is not the only harm done by fake news: there is also the damage done to the institution of journalism as a whole. Since we said that fake news is an attempt to deceive the news consumer into believing that the source of information is a journalistic one, fake news eventually represents a reputational damage for real journalistic sources. It spoils relations between news producers and news consumers, for the latter find themselves duped when they recognize the poor epistemic value of fake news and the violation of the norms. In the long term, this may cause a crisis of legitimacy for real journalistic sources that can lead to the progressive alienation of news consumers, with all the subsequent economic consequences. So reparations of journalist serve also the purpose of thwarting this further predicament.

Another reason to keep false news separated from fake news is one of conceptual necessity. We already know that conservative accounts are vulnerable to the Habgood-Coote's critique of conceptual necessity, because their notion of fake news is too akin to others which already do the conceptual work that "fake news" is supposed to do. "False news" is one of them, together with "misinformation" and "disinformation". Why should we use "fake news" to indicate news with false content when it is enough to simply talk about false news? Many conservative accounts – the hybrid ones – solve this problem by stating that to have fake news we need to satisfy another subjective condition, other than falsehood of content. But most of the time, as we saw in the first chapter (§5), the addition of such condition does more harm than good, because it elicits new counterexamples. Besides, the addition of subjective conditions concerning intentionality of the news producer exposes again the conservative account to the critique of conceptual necessity (think about "disinformation", "bullshit" etc.).

My definition of fake news avoids this problem, for news with false content must not be automatically categorized as fake news (and vice versa). To be fake news, false news must significantly violate norms of journalism, and we cannot take for granted that a false piece of news did it. Sometimes, as bona fide errors prove, perfect adherence to the norms of journalism is not a guarantee for the truth. In conclusion, fake news is not necessarily false or misleading; and false or misleading news is not necessarily fake news. What we might say, instead, is that fake news is at most *typically* false or misleading. The reason is that norms of journalism reflect good epistemic and communicative practices, and a significant violation of them facilitates poor outcomes of this type. Unverified news can easily turn out to be false, and poorly communicated information can easily mislead news consumers.

Concerning other features of account of fake news, someone may wonder whether a piece of news, at some point, may not deserve anymore their status thus becoming fake news, and vice versa. Pepp, Michaelson and Sterken (2019) illustrate something similar when they say that the status of fake news is dynamic. They highlight that news stories can gain or lose their status as "fake news" over time, depending not only on the producer's initial intentions but also on how the story is subsequently disseminated and received. In concrete terms, Pepp et al. argue that a story which was not originally fake news might become so if it is later shared widely by people who treat it as a product of

journalistic practices when it was not in fact produced that way. Conversely, a story that starts out as fake news might eventually lose that status if a recognized journalist or news organization legitimately absorbs it into a genuine process of verification, contextualization, and editing – even if the story remains inaccurate. On their view, "fake news" is not solely determined by the story's content or by the originator's mindset at the moment of production; the broader social uptake and editorial trajectory can transform its standing. But does this dynamism of fake news apply also to my institutional account? Not in the same way.

On one side, it is relatively straightforward to illustrate how rumors or unverified information can acquire the status of news. Suppose a journalist hears a tip about a politician allegedly involved in corruption. Initially, this tip remains merely a piece of information, not yet meriting recognition as news: it has not been cross-checked for reliability and no broader investigation has been performed. If the journalist diligently verifies the source, consults additional documents, interviews relevant witnesses, and presents the findings in a publication that adheres to the accepted norms of journalism, then the underlying tip has effectively acquired the status of news. In this sense, my institutional perspective accommodates a natural and empirically common transition, where once-uncorroborated information comes to be recognized as news precisely by virtue of journalistic scrutiny.

By contrast, it is more difficult to cite convincing examples of once-officially-recognized news losing its institutional status. In principle, a story could be disavowed by the same news organization if it were exposed as systematically violating journalistic standards – perhaps because it was fraudulent from the start or sourced entirely from fabricated testimony. Theoretically, we could say that the piece ceases to be news once it is discovered never to have deserved that status. Yet such reversals rarely occur in practice; corrections or apologies may be issued, but the published artifact is seldom openly reclassified as "no longer news". Consequently, in most real-life scenarios, transitions from rumor or tip to recognized news are more prominent than transitions in the other direction. A piece of information that "graduates" into legitimate news by satisfying institutional norms is far more common than a case in which established news is formally stripped of its status. That is why my definition of fake news is "news that

does not deserve its status"; the status of news remains, although it is undeservedly retained (see §3 above to explore the reasons why, in my view, fake news is still news).

Another feature of my account is that, since there are two different sets of journalistic norms, there are many ways to violate norms of journalism too. This point further differentiates my account from privative ones, which usually treat standards of journalism as solely reducible to norms of news production. Instead, the existence of norms of news publication allows us to consider new phenomena as fake news; phenomena that other accounts of fake news may not be able to capture (as we shall in the next section).

Violations of the norms of news production are quite easy to grasp. For instance, a news item can report that some alleged facts really occurred without supporting this claim with solid and reliable evidence. One of the most infamous examples of Jayson Blair's journalistic misconduct was his article for *The New York Times* about a missing soldier, Private Jessica Lynch (published in April 2003). In the article, Blair described in vivid detail how Lynch's parents, residing in Palestine, West Virginia, were coping with the news of their daughter's capture and injuries during the Iraq War. However, it was later revealed that Blair had never set foot in West Virginia. The details he provided were entirely fabricated or plagiarized from other sources. The inaccuracies in Blair's reporting were eventually exposed when editors began noticing inconsistencies and contradictions in his stories. This particular article, along with others, lacked proper oversight and corroboration by *The New York Times*. The editors trusted Blair's submissions without verifying his travel records or consulting the primary sources he claimed to have used. This is a clear case of fake news, since it is published in plain violation of one of the primary norms of news production: the scrupulous verification of information. Again, if the story reported by *The New York Times* eventually had turned out to be true, the response would have been the same for the reasons exposed above commenting the case of *Russia Today* (p. 137).

There are several other examples of violations of the norms of news production. News items that are severely biased, report information from non-transparent sources, or do not make a synthesis of information, and so on. All of these cases, plus the ones that I did not mention, are potential cases of fake news. One has to remember, however, that the violation must be significant to qualify as fake news, and that we need to assess

this for each piece of news individually. This goes for violations of both sets of norms of journalism, production and publication. The immediate consequence of this maneuver is that we can spot fake news only on an individual basis, looking each time for evidence of a significant violation of norms. So we cannot properly say that Pizzagate is fake news, for instance. The reason is that general assessments like these only refers to the whole story of the so-called Pizzagate and then brand it as false or misleading; it says nothing about how single articles on the same story deviate from good epistemic and communicative practices of journalism.

This final remark leads us directly to the potential flaw of my account, which is its discordance with common use of the term "fake news". Conceiving fake news as news that violates norms of journalism forces us to admit the existence of true fake news and to resist the move of deeming every false news as fake news, as we already said. This feature can result in counterintuitive outcomes, like when I claim that the case of *Russia Today* may be fake news too. The weirdness of such assessment depends on the fact that falseness and fake news are so tight in the ordinary, extra academic conception that seems almost impossible to have the latter without the former. But as I already said in the previous page, there exists a link between fake news and content falseness even in my account, for violations of norms of journalism represent deviations from good epistemic and communicative practices. Therefore fake news *tends to be* more easily false and/or misleading than real news, precisely because it is at odds with the rules that the institution of journalism has established for itself.

In any case, the definitional project I embarked on here is not a descriptive one. My attempt is not to formulate a definition of fake news that diligently tracks intuitions of ordinary speakers around the concept. I did not test my definition with any potential counterexample. In contrast, I am stipulating a concept of fake news that might be useful for the purpose of separating it from other distinct phenomena. The aim is to disentangle fake news from other concepts that might be confused with it, because they address similar, but not identical, phenomena of communication.

4. **The pragmatics of fake news**

Norms of news publication represent a distinct set of norms of journalism. As we know by now, they embody principles of journalistic communication that cannot be

completely reduced to the ones that govern ordinary face-to-face conversations (§2.1). The inclusion of norms of publication within the broader set of journalistic norms has interesting implications for an account of fake news. In fact, reference to norms of news publication opens the door to new ways of falling short of news. Since these ways are rarely captured by existing account of fake news, I will elaborate more on them in this section. Two premises need to be made, though.

Firstly, I cannot offer a definitive, clear-cut list of norms of news publication. The main reason is that those norms are often a matter of conventions that may vary, even substantially, according to the preferences and the personal intuitions of journalists and editors, the genre of the journalistic publication and cultural influences to which the journalistic source is inevitably subjected. Instead of an exhaustive list, a less ambitious approach is to present some cases and then argue why they are likely to be judged as violations of those norms, regardless of the elements of variation cited above. The idea is that there are ways of presenting information to news consumers that, intuitively, no journalistic source should be willing to accept. In §2.1, I hinted at some of them: a sentence construction that leaves too much room to implicature indeterminacy; the use of eye-catching expressions, forward-referencing and other clickbaiting-like techniques; the presence of expressions that serve the purpose of sensationalism through appealing to emotions; an excessive length of the headline due to information abundance (often coupled with a redundant article that does not add any relevant information); the abuse of figurative language. These violations of the norms of publication, plus others, occur fundamentally along two levels. They consist in either the implementation of certain pragmatic strategies or a failure to stick to the style of journalistic writing. In this section I will focus mainly on the first kind of violation – those that concern pragmatic strategies.

Secondly, the observance of norms of news publication applies to written and oral content and, concerning written content, to both headlines and texts. However, in this section I will focus mainly on violations of norms of publication in news headlines. The motivation that drives this choice is that headlines, together with cover images, are the first element to catch the eyes of news consumers, and sometimes it is the only component of a news item that is actually read by them (Gabielkov et al. 2016, Wang et al. 2016). I will take into consideration both real headlines and invented ones,

depending on the case. Also, scholars have recently started to investigate headlines by applying prominent theories of communication in pragmatics to the study of their meanings and communicative goals. This offers me the opportunity to use such insights for the analysis of fake news that exploits the role of headlines. The main account of communication employed so far for the study of headlines in pragmatics is Relevance Theory (Sperber and Wilson 1995).

I believe that the framework of Relevance Theory dovetails with my institutional account of (fake) news. As explained in Chapter 2 (p. 93), attributing news status to a piece of information confers upon it certain expectations of "newsworthiness." By virtue of being designated as news, the item is presumed sufficiently noteworthy, recent, or otherwise significant to warrant public attention. As we shall see soon, this presumption naturally resonates with the central premise of Relevance Theory. People do not read or watch news merely out of habit. They do so because they expect to gain something from it – whether that is understanding current events, learning helpful facts, or seeing a perspective they can use. In Relevance Theory, humans allocate mental effort to what they judge is likely to be "relevant"; that is, anything bringing enough new insight or practical benefit to be worth the time and energy it takes to process it. When something is labeled "news," people presume it must be important or novel enough to justify paying attention.

In this sense, newsworthiness parallels relevance: both denote a promise that "here is content you ought to engage with, as it potentially enriches your understanding or aids decision-making". This synergy explains why the institutional conferral of "news" status (in Searle's sense) so powerfully shapes how content is interpreted. Readers reflexively allot more credibility, attention, and interpretive charity to something flagged as news, implicitly treating it as relevant by default. Consequently, if a producer violates journalistic norms – whether by fabricating facts or crafting misleading headlines – those same heightened assumptions of relevance can be weaponized to mislead consumers. Relevance Theory thus helps reveal how the news label's guarantee of "newsworthiness" can be subverted at the level of actual language use, illustrating how institutional status and pragmatic interpretation reinforce one another.

In the next two sections you find a synthesis of this cognitive framework and a focus on two innovations of RT that are very useful for our analysis of fake news: strength of

implicatures and dynamism of the context. After summing up those notions, I embark in the analysis of headlines of fake news.

4.1 Relevance Theory: staples

Relevance Theory (RT) is a cognitive-pragmatic framework developed by cognitivist scientist Dan Sperber and linguist Deirdre Wilson that refines and overcomes Grice's theory of meaning and implicatures. Unlike Grice's Cooperative Principle and his conversational maxims, which require speakers to follow rules like truthfulness, informativeness, and clarity, RT proposes that human communication revolves around the search for relevance. Accordingly, the fundamental premise of RT is that human cognition and communication are governed by the pursuit of relevance. Within the framework of RT, every utterance creates expectations of relevance that guide hearers in interpreting meaning, eliminating the need for the observance of explicit maxims like Grice's.

RT moves beyond the code model of communication, which assumes messages are encoded, transmitted, and decoded in a linear fashion. Instead, it adopts an inferential model, where speakers provide evidence for their intentions, and hearers infer these intentions based on cognitive processes[63]. Communication is not about encoding and decoding identical messages, but about generating meaning through context and inference. RT thus introduces two core principles:

1. Cognitive Principle of Relevance: Human cognition tends to be geared to the maximization of relevance.
2. Communicative Principle of Relevance: Every ostensive stimulus conveys a presumption of its own optimal relevance.

The Cognitive Principle of Relevance is grounded in the assumption that humans are biologically and psychologically wired to pursue relevance in all aspects of life, not just communication. Relevance, in this context, is a measure of the cognitive benefits an

[63] Grice's model of communication is for the most part inferential too. However, RT incorporates a more pronounced inferential approach due to the way in which RT conceives the explicit level of meaning, what Grice called "what is said". According to relevance theorists, the recovery of the explicit meaning in Grice is based on a version of the code model yet, since Grice believed that what the speaker said coincides with conventional meaning of the terms of the utterance plus disambiguation and reference resolution. RT instead posits that inference has a role in the recovery of the explicit meaning too.

input provides relative to the processing effort required. Humans are constantly bombarded with stimuli, but our cognitive systems are designed to attend to inputs that are worth processing. RT defines relevance in terms of two factors:

- Positive cognitive effects: these are valuable outcomes, such as gaining new information, confirming beliefs, or revising mistaken assumptions. In RT they are usually called "contextual implications": assumptions coming from the elaboration of both an input and the context from which it comes from (neither input nor context alone). For example, on seeing my train arriving (input), I might look at my watch, access my knowledge of the train timetable (context), and derive the contextual implication that my train is late, thus creating a positive cognitive effect.
- Processing effort: this refers to the mental energy required to process an input. The more effort it takes to process an input, the less relevant it is. For instance, a highly technical paper full of jargon may provide useful information but might require considerable cognitive effort to understand, reducing its relevance. Usually the factors that influence processing effort are recency of use, frequency of use, perceptual salience, ease of retrieval from memory, linguistic or logical complexity and size of the context of communication (Barsalou 1987, 1989; Clark 2013).

In practical terms, humans prioritize inputs that strike a balance between maximizing cognitive effects and minimizing processing effort. If processing an input offers significant cognitive benefits with minimal effort, it is considered highly relevant. Conversely, inputs that require excessive effort without providing substantial cognitive effects are deemed less relevant.

Therefore, relevance is a comparative property among stimuli, and not an absolute property of each stimulus. For example, consider the following utterances:

(a) We are serving meat.

(b) We are serving chicken.

Although both are informative, in most contexts (b) is more relevant than (a) because it offers more specific cognitive effects (but not every time; e.g., consider a context in which what is relevant is whether the food is vegetarian or not: we could argue that (a)

is more relevant if we assume that meat, as a more general term, takes less to process than chicken). It answers not only the question of whether meat is being served but also specifies the type of meat. This increases its relevance by providing additional, contextually useful information.

The Communicative Principle of Relevance builds upon the cognitive principle by explaining how relevance is exploited in communication. Every ostensive stimulus is presumed to be optimally relevant, meaning that 1) it is relevant enough to be worth the audience's processing effort, and 2) it is the most relevant one compatible with the communicator's abilities and preferences. But what does "ostensive" mean?

The notion of ostension plays a crucial role in RT. Ostensive communication refers to communicative acts where the speaker overtly signals their intention to communicate. For example, when someone points to an empty glass to indicate they need more water, this is an ostensive act. In contrast, simply placing the glass near someone without drawing attention to it is non-ostensive.

In ostensive-inferential communication, there are two levels of intention:

1. Informative Intention: The speaker's intention to convey information (e.g., "I want a drink").
2. Communicative Intention: The intention to signal to the audience that they are trying to convey this information (e.g., pointing at an empty glass to prompt the inference that they want a refill).

This dual-layered approach allows communicators to achieve optimal relevance. In producing an ostensive stimulus, the speaker guarantees that it is worth the audience's processing effort and encourages the audience to search for relevance. The audience presumes the communication strikes a balance between cognitive effect and processing effort. For example, if someone waves an empty glass at you, the expectation is that this behaviour is optimally relevant, indicating the person would like a refill, without the need for verbal clarification. This notion of optimal relevance allows for flexibility in communication, even in cases of deception or incompetence. In fact, even a deceptive communicator must still make their message appear relevant; otherwise, it would fail to engage the audience's attention.

Comprehension in RT is not a linear process of decoding but an inferential one, guided by relevance. The theory posits that when a person hears an utterance, they do not simply interpret it in a linear or sequential manner (as in the traditional code model of communication), but rather engage in a non-demonstrative inference process, testing possible meanings and interpretations in parallel until their expectations of relevance are satisfied. To understand the process of comprehension we need to explain how the hearer constructs hypothesis about explicit and implicit meanings of the speaker's utterance. The first idea is that hears do this by following a path of least effort. This means the hearer considers the most accessible interpretations first, stopping once they find one that meets their expectation of relevance. Here is the "relevance-theoretic comprehension procedure" as presented by Sperber and Wilson (2004):

a. Follow a path of least effort in computing cognitive effects: Test interpretive hypotheses (disambiguations, reference resolutions, implicatures, etc.) in order of accessibility.
b. Stop when your expectations of relevance are satisfied (or abandoned).

The main idea conveyed by the comprehension procedure is that a speaker who wants their utterance to be as easy as possible to understand should arrange it so that the first interpretation to satisfy the hearer's expectation of relevance is the one the speaker intended to convey. Thus, Sperber and Wilson continue, when a hearer following the path of least effort arrives at an interpretation that satisfies his expectations of relevance, in the absence of contrary evidence, this is the most plausible hypothesis about the speaker's meaning (but notice that this is just an inference process, and it may well be false in concrete situations; nevertheless, it is the most rational route for the hearer according to RT).

For example, if the conversations between the speaker and the hearer revolves around money and finance, the hearer might first interpret an utterance with an ambiguity such as "He forgot to go to the bank" as meaning that the person forgot to visit a financial institution, rather than a river. This is the most accessible and relevant interpretation of the ambiguous term "bank", given that it fits the likely context (financial obligations, such as repaying a debt).

The second element that we need to consider is that there are two level of meaning in interpreting an utterance: the explicitly communicated assumptions and the implicitly

communicated assumptions. Both levels are computed in parallel by hearers, who continuously adjust their hypothesis on the meaning of the utterance until they reach a satisfying degree of relevance, according to their expectations. RT call "explicatures" the explicitly communicated assumptions. Specifically, explicatures are hypothetical developments of the logical forms of an utterance through inferential routes. The logical forms of an utterance are the conceptual imports of lexical terms that constitute the utterance, and they are usually fragmented and incomplete. To complete them, we need to use contextual cues inferentially to develop the logical form, and so to arrive at the explicatures. Usually, to obtain the explicature we need inferences linked to ambiguities, scalar terms, quantifiers, modals, numerals, indexicals, and other pragmatic enrichments (Carston 2002, Clark 2013).

Instead, implicitly communicated assumptions are implicatures. In RT there is no distinction between conventional and conversational implicatures, or between generalised and particularised implicatures. But implicatures in RT can be organised in two other ways: implicated premises or implicated conclusions and stronger or weaker implicatures. Implicated premises are implicitly communicated assumptions that the speaker communicates as a step for the derivation of further implicatures, the implicated conclusions. Implicated premises must not be confused with simple contextual assumptions, which are instead information that the speaker presupposes the hearer already has or can easily retrieve.

Here are two examples from Clark (2013) that illustrate this difference:

I) Ken: Do you fancy a cup of coffee?
Bev: Actually, I'm avoiding caffeine just now.

II) Ken: Do you fancy a cup of coffee?
Bev: Actually, I'm avoiding drinks with tannin in them at the moment.

In (I), we might reasonably suppose that Ken knows in advance that coffee contains caffein. So Bev, in refusing his invitation, is not also giving him the information that coffee contains caffein, but she takes for granted that Ken already knows it. Then, the assumption "coffee contains caffein" in (I) is just a contextual assumption. It is used to derive implicatures just like implicated premises (e.g. the implicature that Bev would not like a cup of coffee right now), but it is not in itself an implicature because it is not really communicated by Bev. It is just a presupposed assumption. In contrast, in (II) Bev

is giving evidence that she believes that coffee contains tannin, perhaps supposing that Ken did not know this information beforehand, since it is not as widely shared encyclopaedic information as that on caffeine. So in this case the assumption "coffee contains tannin" is an implicated premise, a type of implicature, that Ken will use to construe a range of further implicated conclusions (including again the implicature that Bev would not like a cup of coffee right now). The difference between contextual assumptions and implicated premises will be useful for us in distinguishing two separated cases of problematic headlines.

Another innovation of RT is the idea that implicatures can vary in strength. The strongest implicatures are those which need to be recovered in order to reach an interpretation of the utterance that is relevant enough. For instance, both in (I) and (II) Bev communicates very strongly the implicated conclusion that she prefers not to drink coffee in that moment. Without hypothesizing this implicature, Bev's replies become pointless for they would not be relevant enough for Ken to be processed. Descending further into the degrees of strength, we find then implicatures which the hearer is strongly encouraged but not actually forced to recover. For example, in both (I) and (II) Ken may be encouraged to recover the implicature that Bev is not turning his invitation down because she does not like him, or that Bev will not want to drink anything else with either caffeine or tannin. Although not being the strongest, these implicatures are still quite strongly communicated.

But Ken might go on to consider other possible implicatures for which Bev's utterance offers little to no evidence. Maybe, in (I) Ken might wonder why Bev is avoiding caffein and consider various reasons. For example, he might derive the implicature that Bev has having trouble tolerating the effects of caffein, or that caffeine makes her too much wired and nervous, and so on until he may arrive at very far-fetched implicatures that shade off into entirely unintended implications (e.g. that Bev refuses to fund companies that sell drinks with psychoactive drugs). In general the weaker the encouragement, and the wider the range of possible implicatures among which the hearer can choose, the weaker the implicatures (Sperber and Wilson 1995: 199). The encouragement to recover an implicature depends on two factors: 1) how

much evidence the speaker provides that they intend to convey it; 2) the kind of encyclopaedic information around concepts that the hearer can access[64].

Accessibility of encyclopaedic information depends on the context of interpretation of the utterance, whose determination is in turn dynamic for RT. Unlike traditional approaches that treat context of interpretation as fixed or determined in advance, RT posits that contexts are dynamically selected and constructed during the comprehension process. A critical factor in context selection is the organization of memory and the retrieval processes involved. According to RT, chunks of information about concepts are retrieved from memory when their corresponding concepts are activated during utterance comprehension. These chunks of information are arranged hierarchically, allowing the hearer to retrieve contextual assumptions in steps.

I will illustrate this process through an example of an ever-developing exchange made by Sperber and Wilson (1995):

Mary: I'm starting to get hungry. Would you mind preparing dinner tonight?

Peter: I'm tired.

In this scenario, Peter's response would typically be understood as implying that he is too tired to cook, even though this implication is not explicitly stated. This interpretation is derived from a context that includes both the immediate discourse (Peter is tired) and background assumptions (tiredness might prevent someone from cooking).

As the conversation progresses, new assumptions are dynamically added to the context. For example, if Mary responds:

Mary: The dessert is ready. I'll make the main course.

[64] RT assumes that words encode concepts: schemes of representations that are stored in memory and are activated every time the associated word is received as an input from the outside world (e.g. when the speaker utters them). Every stored concept gives access to three kinds of information: lexical, logical and encyclopaedic. Lexical information is about linguistic features of the word, like its pronunciation; logical information is about deductive inferences that may follow from the proposition containing the word under consideration; encyclopaedic information is about objects, events or properties that fall under the concept, which comes from background knowledge and an individual's own experience of the world. The main characteristic of encyclopaedic information is that, by definition, it varies from individual to individual, while logical and linguistic information should be relatively stable across speakers (Clarke 2013: 245).

Peter would understand this response as a continuation of the earlier exchange, interpreting it in a context that includes assumptions about meals (a meal consists of at least a main course and dessert).

In more specific exchanges, such as:

Mary: The dessert is ready. I'll make an osso-buco.

Peter must incorporate further specific assumptions into his context, such as "an osso-buco is a main course". This shows how context is expanded dynamically, adapting to the specific content of the conversation. But access to encylopaedic information can further expand the context. Peter, the hearer, might expand the context to include assumptions about what an osso-buco is (e.g., a veal dish), and his memory may then retrieve associated knowledge (e.g., that osso-buco is time-consuming to prepare), which further shapes the interpretation[65].

The retrieval of these assumptions follows a path of least effort, meaning that the mind retrieves and processes the most relevant and accessible information first. More complex or obscure assumptions may become accessible later, but only if they are necessary to achieve sufficient cognitive effects. In fact, every new expansion of the context entails more processing effort on the part of the hearer, who needs to retrieve more chunks of information from memory. Therefore, context expansion, per se, decreases relevance overall, unless these increments in effort are counterbalanced by even more gains in cognitive effects.

Besides, chunks of information are not all equally accessible. The ease with which certain assumptions are retrieved during context selection is a function of memory accessibility. Sperber and Wilson (1987) proposed that the most frequent or recently used assumptions are more easily accessible, requiring less processing effort to retrieve. This accessibility influences which contexts are selected first during comprehension. In addition, the size and complexity of mental computations play a role in determining processing effort. Smaller chunks of information – such as assumptions that have been frequently or recently used – are easier to retrieve than larger, more complex chunks

[65] Encyclopaedic information does not exhaust all context expansions. RT identifies other two main types of contextual expansions: 1) Assumptions from recent mental processing (for instance, if a person has been discussing a movie, assumptions about the movie plot might be immediately accessible; 2) assumptions from immediate physical environment (such as noticing that it is raining outside and interpreting an utterance accordingly).

that require more cognitive effort. For example, the assumption "Peter is tired" might be easily retrieved if it has been mentioned previously in the conversation. In contrast, recalling the assumption "Peter often cooks when he's not tired" might require more cognitive effort because it involves retrieving multiple related assumptions.

Other scholars propose further mechanisms that might influence the accessibility of chunks of encyclopaedic information. Here is one thesis from Assimakopoulos (2017):

> I think it would be uncontroversial to claim that the more some piece of information has in its processing history affected our belief system, the easier it should be to retrieve in future processing, or in RT terms, the more positive cognitive effects an assumption has helped produce in a previous inference, the more accessible it should get for future inferences too and, thus, the more frequently it would be used in them. (Assimakopoulos 2017: 231)

Assimakopoulos argues that assumptions that have produced significant cognitive effects in the past are more easily accessible in the future, even if they have not been recently activated. This proposal builds on the idea that relevance is not just about the immediate context or recent assumptions but also includes past assumptions that have had a profound impact on an individual's system of beliefs.

The idea that past relevance enhances future accessibility means that a hearer does not need to start from scratch when selecting contexts for interpreting new utterances. Instead, the mind can automatically retrieve relevant assumptions from the past based on their previous impact. For example, if someone once had a near-death experience in a car accident, the assumption "Driving can be dangerous" may not be actively processed every time they get into a car. However, when someone mentions a new car accident, this deeply relevant assumption can be immediately accessed to interpret the conversation, without significant cognitive effort.

In practice, this means that assumptions that have played a central role in an individual's system of beliefs – whether from recent or distant past – are pre-selected during comprehension, even if they are not frequently activated. This process would occur without added cognitive effort because these assumptions are already primed for retrieval due to their cognitive effects.

The plan for the next sections is to apply RT-based analyses of headlines to identify what is wrong in the headlines of some fake news. Specifically, I will use the notions of

RT summarized here to describe example of headlines which are significantly at odds with the norms of news publication. In doing so, I will identify three classes of problematic headlines: clickbait headlines, deceptive headlines and targeted headlines (or targeted fake news).

The first class is a well-known type of headlines crafted according to the strategy of forward-referencing. Forward-referencing involves the use of incomplete and/or vague phrases that serves the purpose of pushing the reader towards the landing site by exploiting an information gap. In developing my analysis of clickbait headlines, I will follow the insights of Scott (2021), who has proposed a RT-based explanation of forward-referencing.

The second class of headlines, instead, was recently coined by Jodłowiec (2023). Like clickbait headlines, deceptive headlines are designed to intrigue the reader and to push them to the landing site. But deceptive headlines do not rely on forward-referencing; rather, they resemble more standard headlines that provide a summary or an overview of the accompanying text.

Finally, for targeted headlines I mean a new class of headlines that exploit weak implicatures to deliver messages to specific audiences. In my view, they represent a subtle way to violate norms of news publication, for target headlines are formally true and unproblematic at first glance. Target headlines will be extremely important for our case study on fake news about climate change in the next chapter.

4.*2 Clickbait headlines*

Clickbait headlines refer to headlines or links that are deliberately designed to provoke curiosity and drive online traffic by encouraging users to click. These headlines are often sensationalist or misleading, promising content that will reveal surprising, shocking, or important information. Despite the high expectations set by the headline, the actual content often fails to deliver, leading to user dissatisfaction. However, from a marketing perspective, clickbait is effective because the goal is to generate a click, which leads to ad revenue or engagement, not necessarily to satisfy the reader's informational needs.

Many conservative accounts of fake news defend a subjective condition to account for what is wrong in clickbait headlines. As we saw in the first chapter (p. 31), most

hybrid accounts count clickbait headlines as fake news because they maintain that the author of the headline is driven by bullshitting intentions; but this thesis opens the door to many objections (first chapter, §5.2). In contrast, my institutional account of fake news locates the wrong of clickbait in the violations of the norms of publication. The main idea is that clickbait headlines deviate from the stylistic features of journalistic writing, for they *distort the objective of news headlines*. But what is this objective?

Recent works support the view that headlines must be seen as *autonomous texts* rather than summaries or anticipations of the story they introduce (León 1997; Smith 1999; Althaus et al. 2001; Andrew 2007, Ifantidou 2009). This autonomy means that the headline can be engaging, intriguing, or meaningful on its own, without the need for readers to access the entire article. News consumers may access to the article in order to find a full picture of the story, but this remains just a possibility: the consumption of a news headline remains always independent and autonomous from the consumption of the connected article. But this is exactly the opposite for clickbait headlines.

Rather than being self-contained, autonomous texts that can be purposely read for the sake of a quick and loose news update (Ifantidou 2009: 702), clickbait headlines most of the time are forward-referencing: they build a bridge between them and the article, so that consumption of the latter becomes a necessary requirement for the cognitive processing of the former. The bridge is formed through the exploitation of an information gap: the headlines is pointing to something that cannot be fully grasped without reading the text. This means that clickbait headlines cannot be autonomous texts. Consider for instance these two invented headlines from Koechley (2012):

(3) Obama says gay marriage should be legal.

(3*) Now THIS is why I voted for Obama.

(3) is an autonomous text inasmuch as it does not need anything more than itself to be correctly processed and understood by the reader. If the news consumer wants to delve more into the news story, they can open the link and read the whole article; but this action is neither necessary nor required for the comprehension of the headline. On the contrary (3*), which is the clickbait version of (3), is not a self-contained headline. News consumers do not satisfy their cognitive needs through the reading of (3*): they must also read the accompanying text. Scott (2021) has argued that Relevance Theory

provides an ideal framework to understand how clickbait headlines exploit cognitive mechanisms to entice readers into clicking the link to read the text. Thus, we can explain the difference between (3) and (3*) through the lens of RT, following the argument and the analysis made in Scott (2020).

To determine the explicature of (3), the reader only needs to identify the referent of the name "Obama". Reference resolution is a crucial part of the determination of the explicit content, involving mapping a mental representation of the intended referent onto the logical form of the utterance (Scott 2020). In the case of "Obama" in (3), there is likely one highly accessible mental representation of the referent: Barack Obama, who was the President of the United States at the time of the publication of (3) in Koechley (2012). So the reader of (3) typically resolves this reference by associating the name "Obama" with their existing mental file for Barack Obama.

This process results in a truth-evaluable proposition (3E), which is the explicature of (3) (adjustments are in square brackets):

(3E) [Barack] Obama [, the current president of United States,] says gay marriage should be legal.[66]

Notice that this explicature can be further extended through other pragmatic enrichments, as in (3E'):

(3E') [Barack] Obama [, the current president of United States,] [recently said that] gay marriage should be legal [within the territory of United States].

But most of the time versions like (3E) are more than sufficient for the expectations of relevance of the reader. That is, it is most likely that (3E) will achieve optimal relevance, and the reader will accept it as the intended interpretation of (3).

Taking (3E) as the correct interpretation of (3) can lead to various cognitive effects. For example, Scott suggests that it might strengthen the reader's existing beliefs about Barack Obama or interact with other assumptions they hold about him, potentially leading to new insights about him or related topics like gay marriage. Consequently, the reader updates their conceptual file for Barack Obama with these new assumptions: now the assumptions encoded by (3E) is stored in the concept of "Barack Obama", and every

[66] I am assuming the perspective of someone who reads (3) from Koechley (2012): at that time Barack Obama was the President of United States.

time this concept is at stake the speaker might retrieve this assumption (more or less strongly, depending on the context).

Now we turn to (3*), the clickbait version. In this case, to calculate an explicature, Scott (2020) says that the reader must find the reference for three expressions: "Obama", "THIS" and "I". While resolution of the first element is likely to be the same of (3), there are no obvious candidate referents for the other two expressions. If left unresolved, the headline offers only two pieces of information to the reader: 1) there is a person who voted for Barack Obama; 2) that person voted him for a specific reason. But these two assumptions alone cannot give place to many cognitive effects: readers usually assume that there are people who voted for Barack Obama and who did so for a particular reason. It is very poor information.

But as the Communicative Principle of Relevance says, every ostensive stimulus conveys a presumption of its own optimal relevance. Therefore, the entertainment of these assumptions must be justified somehow, for the reader assumes that the writer thought that those assumptions would be optimally relevant for them. This prompts the reader to try to work out how those assumptions might be optimally relevant. And here comes the role of demonstratives like "this", which are one of the most widespread linguistic markers of clickbait headlines[67].

In fact, demonstratives are used to point something or someone out to the reader. By capitalizing "this", the author of (3*) signals that it is significant. Drawing attention to something is an intentional act that creates an expectation of relevance, whether done through a physical gesture or a linguistic marker. In this case, using "this" in the headline is akin to pointing toward a specific piece of information, leading the reader to reasonably expect they will discover what "this" represents by following the link and reading the accompanying article. In doing so, the author explicitly directs the reader's focus to the reason behind the vote for Obama, which are still unknown. In other words, the use of the demonstrative in (3*) raises our expectations that following the link will be optimally relevant. It is a promise that the reader will achieve cognitive effects that

[67] Other linguistic markers of clickbait headlines are superlatives and hyperbole ("You Won't Believe the Best Way to Lose Weight!"), ellipsis and open-ended phrasing ("You'll Never Guess What Happened Next..."), second-person pronouns ("Why You Need to Start Eating This Superfood Right Now!"), numerals and lists ("10 Tips You Need to Know Before Traveling!"), and many others. On the same topic, see for instance Blom and Hansen (2015), Chakraborty et al. (2016).

live up to the expectations raised by "this" only by reaching the landing site. But articles that accompany clickbait headlines rarely respect this promise.

The takeaway of this analysis is that we do not need to dig into the intentions of news producer to spot the wrong in clickbait headlines. News with clickbait headlines can potentially be fake news because the practice of clickbaiting severely distorts the nature and the role of journalistic headlines, which are supposed to be autonomous and self-contained texts[68]. Instead, clickbait headlines are structurally dependent on the accompanying article because they are not optimally relevant on their own right; their sole consumption does not repay the processing effort with as many positive cognitive effects.

We can spot clickbait headlines, and so violation of norms of news publication, by noticing whether the headline in question adopts one or more of the linguistic markers of clickbaiting (footnote 67, above). Usually, it is quite easy to do, and the many the markers, the more significant is the violation. Consider the following headline taken from an online sample of false news gathered by *BuzzFeed News* during 2016 US election campaign:

(4) WikiLeaks CONFIRMS Hillary Sold Weapons to ISIS… Then Drops Another BOMBSHELL! (Silverman 2016a).

The fact that (4) contains false information – Hillary Clinton never sold weapons to ISIS – does not interest us here. What I would like to point out is that (4), like (3*), fall short of our expectations of journalistic writing, and that is why it deserves to be dubbed as fake news. Here we find three linguist markers for clickbaiting: the use of capital letters with the intent to emphasize something, the use of eye-catching words ("bombshell") to arouse curiosity and grab the attention of the reader, the deliberate ellipsis of the second alleged 'bombshell'. These linguist markers serve again the

[68] Again, I say "potentially" because the concept of fake news points to a continuu, (pp- 125-126, above). Situating fake news along a continuum helps clarify why certain headlines – though sensational and misleading – remain recognized as subpar journalism rather than outright disqualifying their status as news. At the same time, it explains how news-like stories, subject to serious *or* systematic violations, lose their rightful claim to being news and become properly labeled as fake news. This perspective captures the dynamic, evolving nature of modern media practices, where occasionally credible outlets can still engage in sensational clickbait strategies, and formerly reputable authors might lapse into repeated or egregious distortions – suddenly pushing their content across the line into fake news. Henceforth, though, I will talk about clickbait as if it is enough to determine plain violation of norms of journalism for the sake of simplicity.

purpose of inducing the reader into clicking the link to satisfy their cognitive needs, thus making the headline (4) a non-autonomous text.

However, I would like to emphasize that, unlike Scott and other authors, I do not consider "forward-referencing" in headlines solely as a linguistic or explicitly future-oriented device (as in the cases just presented), but rather as any structural choice that deliberately withholds key information so that the reader must progress into the story to uncover it. Specifically, the hallmark of clickbait rests on curating an information gap – a point of curiosity or suspense – such that the missing detail is accessible only by 'looking ahead'. In practice, clickbait headlines achieve this also through varied rhetorical strategies: they may tease a surprising revelation without stating it outright, present a puzzle-like scenario whose resolution lies further on, or invoke an emotionally charged premise that promises clarity or closure in the main body. All these techniques manifest the same fundamental mechanism: they extend an implicit invitation for the reader to continue, promising that some crucial explanatory or sensational element awaits. Therefore, in my view, clickbait is a headline practice whose essential property is that it requires the reader to move forward – by clicking, tapping, or otherwise continuing – to discover the core piece of information the headline withholds or only teases. All rhetorical or emotional triggers (curiosity gap, puzzle, etc.) function by referencing that future portion of content. By casting a wider net around what we call forward-referencing, we can incorporate both "classic" clickbait phrasing – where the headline explicitly hints at impending shock or excitement – and less overt but equally withholding styles of headline writing.

4.3 Deceptive headlines

Deceptive headlines refer to a class of headlines firstly diagnosed by Jodłowiec (2023). Jodłowiec defines deceptive headlines as manipulative headlines whose meaning is incongruent with the content of the linked article. While not overtly false, deceptive headlines exploit semantic indeterminacy, leading the reader to form an incorrect interpretation. As in the case of clickbait headlines, deceptive headlines lure readers into clicking by arousing their curiosity, only for them to find that the article does not match their expectations. However, contrary to the clickbait class, deceptive headlines are not

forward-referencing: they are autonomous texts like normal journalistic headlines[69]. So, what do they do to violate norms of news publication? Examples may help us.

A widespread group of deceptive headlines described by Jodłowiec are the ones adopted for death hoaxes. A death hoax is the deliberate presentation of a report about the death of a person that actually never occurred. Usually, death hoaxes target celebrities and well-known people, for their death would be interesting for a much wider audience. Headlines adopted for death hoaxes range from blatant lies (e.g. "the person x has died" when x is alive) to more sophisticated forms of deception. Jodłowiec's deceptive headlines belong to this second category. In her paper she also offers an example of a deceptive headline that appeared for a while on a Polish news website:

(5) The Polish national team coach dies.

In the case of (3), to interpret (5) the reader must assign reference to the noun phrase "The Polish national team coach". For the general Polish audience, the most accessible interpretation of the refence is the coach of the current Polish national men's soccer team. Such an interpretation would lead to a great batch of cognitive effects for that audience, because the death of such a high-profile individual would be highly relevant to many readers. This interpretation would be also in line with the idea that news items are expected to be newsworthy, that is worthy of our attention as news consumers. Journalists should select information that is potentially interesting for us, either because it is significatively impactful for our lives or it arouses our curiosity. The interpretation of the noun phrase in (5) as "The current Polish national men's soccer team" aligns with our general expectations of newsworthiness.

However, when readers click through the link attached to (5), they find that their interpretation of the reference of the noun phrase was wrong. The Polish national coach referred to in (5) is not the national men's soccer team coach. Instead, it is the coach of the Polish junior volleyball team, who is much less well known than the one envisaged by readers of (5). So expectations of news consumers get eventually frustrated. Notice that contrary to many cases of classic clickbait headlines, the article that accompanies (5) is still newsworthy to a certain degree, for the death of the junior volleyball team

[69] It is important to note that, unlike me, both Scott and Jodłowiec consider deceptive headlines as a subclass of clickbait headlines.

may still be upsetting for people who follow that team. The problem is that the article is not as newsworthy as its headline promised to be, not even remotely.

So the first reason of why deceptive headlines deserve the label of "fake news" is that they fall short of our expectations of newsworthiness. They carry the promise that readers will find information worthy of their attention as long as they trust their first, more plausible interpretation and click through the link. But this not the only wrong in deceptive headlines. There is a second one that has to do with the communicative responsibility of the author who craft those headlines. To see what it is let's consider the following fictious headline:

(6) The world will miss Celin Dion's music forever.

Suppose that (6) accompanies and article that reports that Célin Dion is ill and has indefinitely interrupted her career due to her health condition. Even though the explicit interpretation of (6) is true and aligns with the content of the article – Célin Dion really stopped her career for medical issues – the way in which the headline presents the fact is potentially misleading, for it carries the *strong implicature* that Célin Dion is dead. Here is why from an RT perspective.

We know that an utterance is optimally relevant when it is relevant enough to be worth the audience's processing effort (p. 148), and it is the most relevant one compatible with the communicator's abilities and preferences. In the case of (6), the interpretation that makes it optimally relevant is the one that include the implicated conclusion that Célin Dion is dead. The reason is that this interpretation is relevant enough to be worth processing, since it is highly newsworthy (i.e. it leads to great cognitive effects). And if the readers have no reason to believe, as it is normally the case, that the author of (6) was in some way unable or unwilling to produce an alternative, more relevant utterance, then they can reasonably accept "Célin Dion is dead" as the implicated conclusion of (6). Here is how we might analyze (6) according to the relevance-theoretic comprehension procedure.

Firstly, by interpreting "the world" as a figurative speech and enriching the utterance with temporal reference we may obtain (6E) as the possible explicature of (6):

(6E) [Henceforth] [every person who has at least once listened to Célin Dion's music] will miss Célin Dion's music forever.

Then we may recover two related implicit premises of (6):

(6IP$_1$) In order to miss someone's music, there must be a reason that forced them to stop producing music forever.

(6IP$_2$) Death is a clear reason of why someone would stop producing music forever.

From the combination of (6E), (6IP$_1$) and (6IP$_2$) we might recover the already mentioned implicated conclusion that Célin Dion is dead. This conclusion is strongly implicated because its recovery is needed in order to arrive at an interpretation that satisfies the readers expectations of optimal relevance. So the author of (6) encourages the news consumer to derive this implicated conclusion.

This assessment is further supported by the fact that if author of (6) was not suggesting that proposition, there were various alternative headlines that they could easily have chosen to avoid misunderstandings. For instance, they could have chosen the following alternative headline:

(6*) An unexpected illness stops Celin Dion's career indefinitely.

And the same goes for (5):

(5*) The national junior volleyball team coach dies.

Following Scott (2023), the point is that the authors of (5) and (6) should be aware of how these headlines would be interpreted, and so if they genuinely intended to accurately represent the content of the accompanying articles, they should have adapted their headlines accordingly. But the authors of (5) and (6) did not do this. Instead, they let the readers derive either an incorrect explicature, as in example (5), or a false implicated conclusion, as in the case of (6). In both cases the authors of these headlines have violated norms of news publication inasmuch as they did not make the correct proposition easily accessible for the news consumer. In fact, as we said in §2.1, journalistic communication does not take hold in a context of mutual adjustment. Therefore, journalists have a pressing duty in making their intended messages as accessible as possible for the readers. And that is the second reason of why both (5) and (6) are fake news in my account: as deceptive headlines, they are produced without taking into consideration this duty.

It is important to note that my assessment of deceptive headlines does not depend on the truth-value of the propositions they convey. I am not claiming that, say, (6) is fake news because the implicated conclusion "Célin Dion is dead" is false. To avoid this confusion, consider another fictitious headline:

(7) Sputnik V, latest data on vaccine efficacy eventually uncovers the truth.

Imagine that the author of (7) knows that there was an ongoing debate on the efficacy of the vaccine Sputnik V, and that those latest data put an end on it. So the explicit meaning of (7) is true. However, (7) gives rise to a range of possible implicatures that may be ordered on a continuum that goes from the strongest ones to the weakest ones (1= strongest; 4 = weakest):

(IP_1) There was an unsettled debate over the efficacy of Sputnik V.

(IP_2) Telling the truth on the matter took quite some time and effort.

(IP_3) The latest data on Sputnik V were unexpected.

(IP_4) There was an attempt to conceal the truth on Sputnik V's efficacy.

Suppose that each implicature is true and the titlist who wrote (7) intended to convey only (IP_3) as the main implicature of their headline. This becomes clear when the news consumer reads the accompanying article. Although (IP_3) is true, (7) still counts as fake news in my view. The reason is that it leaves too much room for implicature indeterminacy. Although (7)'s calculable implicatures are true, including (IMP_3), the headline does not make the intended one immediately accessible to the readers, leaving them the whole responsibility for the interpretative process. If the author of (7) wanted to convey only (IP_3), they should have written the headline differently. For instance, they could write "Sputnik V, the latest data are at odds with previous ones" or the like. Thus, violation of norms of news publication does not necessarily entail falsehood of content, and as in the case of unverified but true news we can have cases of fake news with true explicit/implicit content.

To recap, deceptive headlines are fake news for two reasons. Firstly, they fall short of our expectations of newsworthiness, because they prompt in the reader an inflated interpretation of what they are getting across, which is usually less important and interesting than what readers envisaged. Furthermore, authors of deceptive headlines

violate the norms of news publication because they do not make the correct explicit or implicit message as accessible as possible for the reader; thus falling short of their obligations as communicators in a context that lacks mutual adjustment.

4.4 Targeted fake news

Targeted fake news is a brand-new kind of news that violates norms of journalism in a strategic manner. They represent a peculiar way to fall short of realistic expectations of news consumers towards journalistic products, since they exploit the difference between the *audience* and the *target* of news.

In a recent book, Mercier (2020) brings back the story of one of the most bizarre frauds of the early 2000s to cause the loss of thousands of dollars all over the western world: the Nigerian scam. Still currently in circulation, this long-running fraud goes like this: someone overseas who claims to be Nigerian royalty emails you offering a share of a huge investment opportunity or a fortune he cannot get out of the country without your help. Then he asks you to send an advance payment to allow him to cover the expenses of the money transfer. As soon as the victim makes the payment, the man disappears, leaving the mark empty-handed (in the worst cases, he even drains the victim's bank account). After recounting the fraud, Mercier poses the question how such a ludicrous scam can be even remotely successful. Are ordinary people so gullible?

Mercier's answer is interesting for our purposes. He begins by noting that for every person who gets played by the Nigerian scam, there are indeed thousands of people who simply ignore or laugh at emails of this kind. That is, only a tiny portion of the whole audience receiving the deceptive email ends up accepting the scammer's proposal. As some may notice, this is the same observation that led to our distinction between consumers and receivers of fake news in the first chapter (p. 37). But here the role of the ridiculous nature of the scam kicks in. According to Mercier, the content design was driven by economic calculations: while sending thousands of fraudulent messages is practically free, not so responding to all of them. Therefore, if the scammers had designed the fraud in a more credible way, they would have ended up with many more potential victims who would have had to be responded to one by one, each in their own thread. This would have entailed far greater effort on the part of the scammers in terms of tailored replies. In contrast, designing a ludicrous scam greatly diminishes the chance

of getting replies from intended victims, thus at the same tremendously reducing the efforts required of scammers. The reason is that only very gullible subjects would be apt to reply to such an absurd proposal as the one from the bogus nobleman. This means that the benefits of the rare successes always exceed the costs of the scam, which ultimately boil down to sending preposterous emails haphazardly.

In my view, the takeaway from the Nigerian scam is that we should always distinguish the *audience* of a message from its *target*. The former is simply the group of consumers reached by a single message. The size of that group may be large or small, but that makes little difference. The same distinction applies to news in general. As in the case of the Nigerian scam, the audience of news usually includes subjects who disbelieve or ignore the news because it strikes them as not credible or simply irrelevant to their interests (recall that consumers are those who read news without necessarily accepting it in any sense). If producers of news wanted to have an effect on the audience at large, such as generating beliefs or prompting acts of sharing, they should take into account the whole audience's many subjective interests, then average them somehow and eventually spread the final message accordingly. But producers do not usually reason in this way: they do not take into account the whole potential audience of their news. That would mean committing the same mistake as 'Nigerian' scammers deciding to make their fraud more credible. Instead, producers may rely on a specific *target*. The news target is often a subset of the whole audience: the subset whose interests and beliefs best match the ones made salient by the news item in question. While 'Nigerian' scammers design their own fraud aiming only at the most gullible victims, producers of news likewise may design their own news content to attract only specific portions of their audience – namely, the ones most prone to accept it. News producers sometimes aim at their own target group of probable receivers[70]. They are aiming at concentrated circulation of their journalistic products.

Crafting and disseminating a piece of news with a target in mind is not a bad thing per se. News producers who are publishing news about a regional event, such as the election of a town mayor, are disseminating news that is primarily directed at a specific

[70] The difference between audience and target is not in number of subjects, but rather in beliefs and preferences: the target group is one characterised by more homogeneous preferences than the general audience, whose interests are heterogeneous by definition. This clarification should help us understand why sometimes the target group seems to be large in absolute terms: we must always compare it to the general audience in relative terms, acknowledging their differences with regard to beliefs and preferences.

subset of the whole potential public (viz. the citizens of that town). But in doing so they are not violating norms of journalism in any way. Most of the time the target of news is also part and parcel of the foundation of a journalistic source, as in the case of local newspapers. So targeted news is not necessarily fake news. It becomes fake news when the target group of news consumers is *covert*. But how to discriminate between overt and covert target groups? Here you find my proposal.

A target group is overt when non-target consumers, although unable to achieve the same degree of cognitive effects of target consumers, are still capable to recover the correct inferential process and achieve a great portion of those effects. For instance, when a local newspaper, dealing with the election in the town Y, reports that "X has been elected the town mayor", people who are not affected by this information, maybe because they live elsewhere, can still derive the intended explicature and, if any, the accompanying strong implicatures; even though they cannot derive all the contextual implications of the inhabitants of Y. For example, they cannot update their set of beliefs with the following one: "X is now the mayor of my town" (and all the ones that may come from it as consequences).

In contrast, a target group is covert when non-target consumers cannot recover the whole inferential process, and therefore achieve a significantly lower degree of cognitive effects than the one achieved by target consumers. One way to achieve this outcome is through the exploitation of different sets of encyclopaedic knowledge around the same subject (but see the next chapter on climate change for other strategies). We can find an example of this strategy in many news headlines appeared in Italy during the anti-Covid vaccination campaign in 2021/2022.

At that time, anti-vaccine blogs were emphasizing the very rare side effects of anti-Covid vaccines to fuel doubts among population. To pursue their goal, a prominent tactic on social media was to re-share compulsively any local news about sudden deaths, which were almost always causally unrelated to vaccination, to artificially create a link between those casualties and the anti-Covid vaccine. Very soon, many local Italian newspapers saw a sudden increment in views and comments under their obituaries: people who hold strong anti-vaccine beliefs were flooding the comment section with more or less explicit references to m-RNA vaccines, blaming them for the sudden death in question. Due to the increment in interactions, this behaviour stimulated

a greater demand of obituaries on the part of local newspapers, which, sometimes cynically, started to report cases of sudden deaths with much more frequency. This, in turn, gave rise to a self-fulfilling prophecy: newspapers were focusing more on obituaries, then anti-vaccinationists found more articles about sudden deaths, which fuelled their misconceptions about vaccines and prompted them to comment more, thus demanding even more obituaries from newspapers, and so on and so forth.

The vicious cycle was also reinforced by a blameworthy act of omission: local newspapers, which sought that enormous wave of misguided claims in their comment sections, never took measures. They neither issued a correction to their articles to clearly state that there was no evidence linking the vaccine to the casualties, nor did they change the way they communicated those facts. Indeed, many of them took advantage of the situation to increase social engagement of the obituaries. This omission contributed to reinforce the mistaken beliefs of the anti-vaccinationists.

From the perspective of newspapers that were publishing obituaries about sudden deaths, the anti-vaccine groups represented the covert target of their articles. The prototypical headline for such obituaries used the expression "malore improvviso" ('sudden illness') to talk about stories of a sudden death. So anti-vaccinationists learned with time that the same expression should be associated with side effects of vaccination. Below an example with a fictitious name:

(8) *Malore improvviso: Caterina muore a 25 anni.*
Sudden illness: Caterina dies at 25

Headlines like (8) were often accompanied by a text that stresses how unexpected the death was. For instance they may have reported that Caterina was thought to be completely healthy, or they may have highlighted the fact that she was very young. Notice that most of the time these articles were not newsworthy at all. Contrary to death hoaxes, in fact, people who appeared in those obituaries were not well-known personalities, and so their demise, as ill-fated as it was, could not achieve per se a great number of cognitive effects in the readers (apart from the deceased's relatives and acquaintances). So how did headlines like (8) and their connected text achieve relevance for their target? By leveraging on three factors: the mutual cognitive environment between the writer of (8) and the target group, and the specific, anti-vaccine

encyclopaedic information around the expression “sudden illness”. Let’s see all of them in order.

Within Relevance Theory, the cognitive environment of an individual is the set of assumptions that are manifest to them at a given time (Sperber and Wilson 1995: 39). An assumption is manifest to an individual at a given time when they are capable at that time of representing it mentally and accepting it as true or probably true. So assumptions are manifest for an individual when they are potentially perceptible, like sensory data, or inferable, like the implicatures of an utterance. For instance, the assumption that my computer is now turned on is manifest right now for me writing these lines of my doctoral thesis (but it may not be manifest any more when you will read this paragraph)[71].

Besides, the cognitive environments of two (or more) individuals can partially or totally intersect at a given time. The total shared cognitive environment of two (or more) people is the set of all assumptions that are manifest to them both. For instance, the assumption that my computer is now turned on can be manifest in this moment also for the deskmate that sits next to me, right now, in the library. Every time the assumption that two (or more) people share a cognitive environment is manifest for both individuals we have a mutual cognitive environment. If both my deskmate and I can potentially accept as true the fact that the assumption “my/his computer is turned on” is shared by both of us, then we also share a mutual cognitive environment.

A mutual cognitive environment is then a cognitive environment which is shared by a group of individuals and in which it is manifest to those individuals that they share it with each other (Carston 2002). Every manifest assumption in a mutual cognitive environment is said to be “mutually manifest”. Ostensive-inferential communication, for RT, is just the attempt of a communicator to produce a stimulus (e.g. an utterance) which makes it mutually manifest to the communicator and their audience that the communicator intends, by means of this stimulus, to make manifest or more manifest to the audience a certain number of assumptions (Sperber and Wilson 1995: 63).

[71] Notice that manifestness is weaker than knowledge. An assumption can be manifest and yet be false, whereas I cannot really know a false assumption, for knowledge entails veridicality. Also, an assumption can be manifest without being entertained by the subject. It is sufficient that the individual is capable to do so; that is why we said that a manifest assumption is perceptible or inferable (and not perceived or inferred).

My starting hypothesis for the analysis of (8) is that the titlist and the news consumers, be them target or not, share a set of assumptions about the SARS-CoV-2 pandemic and the related vaccination campaign. For instance, we may reasonably suppose that assumptions like "There is an ongoing SARS-CoV-2 pandemic", "Young people are vaccinating against SARS-CoV-2", "If any vaccine has side effects, then vaccines for Covid-19 have side effects too", and many others were all shared by the titlist and the news consumers. Therefore, the author of (8) could reasonably assume that an assumption like "Vaccines for Covid-19 have side effects" is shared by people who read (8). And the same goes the other way: readers of (8) could reasonably assume that the author of (8) shared the same assumption (or the like). Therefore, we can say that the assumptions

(8A) *La gente si sta vaccinando contro la Covid-19.*
People are getting vaccinated against Covid-19

(8B) *I vaccini per la Covid-19 hanno effetti collaterali.*
Vaccines for Covid-19 have side effects

were mutually manifest assumptions for both the author and all readers of (8).

Now considers Italian anti-vaccinationists, the covert target group of (8). For them, (8A) and (8B) are manifest too. But as anti-vaccinationists they usually entertain other stronger assumptions about vaccination that, as soon as they read (8), suddenly become *more manifest* to them; for instance, "Vaccines for Covid-19 are lethal for young people" and the like[72]. This is due to the encyclopaedic import of the expression "malore improvviso" ('sudden illness') among the target group, which determines a context expansion for them. To appreciate this effect, let's compare how target consumers and non-target consumers would process an utterance like (8).

Firstly, a target consumer who did not know the deceased will recover the explicature (8E) out of (8) by reordering the proposition and resolving the reference:

(8E) *Caterina, una ragazza non nota, muore a 25 anni a causa di un malore improvviso.*
Caterina, an unknown girl, dies at 25 as a result of a sudden illness

[72] I speculate this assumption based on a review of both the literature about Covid-19 disinformation and the fact-checking articles written during the vaccination campaign. As an example, check an article of Reuters Fact Check (2022) that deals precisely with disinformation around vaccination of young people.

We can discuss whether we should treat the suggested causal connection between the two events as part of the explicature or as a very strong implicature of (8). In my view, colon (":") in (8) can follow the same treatment of "and" in relevance theoretic accounts, as in Carston (2002). The causal and temporal interpretation of "and" in some utterances is taken to be part of explicature in RT-based accounts, for temporal and causal dimensions affect the truth-conditions of utterances. Since ":" in (8) performs the same function – that is, signaling both a temporal and a causal interpretation between the two events (the onset of a sudden illness and the decease) – I shall follow the orthodox treatment of "and" in RT and consider the causal connection as part and parcel of the explicature (8E).

(8E) alone can satisfy the expectation of relevance of non-target consumers only until a certain degree, which is presumably very low. In fact, if they did not know Caterina, they can at most get the information that an unknown person has died, which is not so newsworthy. We can imagine that non-target consumers may then entertain a range of possible weak implicatures of (8), like "Caterina's death was unexpected due to her age", or "Caterina's family is devastated" (both potentially cancellable in the text). But this is hardly enough to prompt them to read the accompanying article or to interact with it in other ways (e.g. commenting under the social media post of the local newspaper). Therefore, non-target consumers will probably brand the article as uninteresting for them and disregard it.

It is a very different story with target consumers. Italian anti-vaccinationists have in fact learned, through repetition, to interpret the expression "malore improvviso" ('sudden illness') as denotating the occurrence of a deadly side effect of Covid vaccines[73]. In RT-based terms, their concept encoded by "malore improvviso" stores, like a mental file, the encyclopaedic information that a sudden illness is usually the deadly outcome of anti-Covid vaccination. This peculiar (and wrong) chunk of information may not be immediately accessible to the target consumer, for we said that accessibility of encyclopaedic information attached to concepts depends on the context of interpretation of the utterance (p. 152). However, communication between the titlist and the target consumers of (8) occurs in a mutual cognitive environment, for we have

[73] There is evidence that an analogous journalistic expression – "died suddenly" – was associated with the same outcome among English-speaking anti-vaccinationists (Goldhamer 2022), thus potentially replicating the same phenomenon outside Italian 'borders'.

assumed that at least (8A) and (8B) are mutually manifest for both the author and all readers of (8). Therefore, contrary to non-target consumers, anti-vaccinationists can fruitfully utilize (8A) and (8B) for the comprehension procedure of (8). In fact, the mutual manifestness of (8A) and (8B) makes that peculiar chunk of encyclopaedic information about "malore improvviso" more easily accessible for them, thus expanding the context of (8) and developing a range of further cognitive effects.

From the combination of (8A), (8B), (8E) and the peculiar encyclopaedic information attached to the concept of "malore improvviso", the target consumer may then retrieve (8IC) as the primary implicated conclusion of (8):

(8IC) *Caterina è morta a causa di un effetto collaterale del vaccino per la Covid-19*.
Caterina died due to a side effect of the Covid-19 vaccine

Notice two things. Firstly, even though (8A) and (8B) were mutually manifest also for non-target consumers, the latter group of readers cannot derive (8IC) from the elaboration of (8) alone, or it is at least very unlikely. At most, some non-target consumer might derive (8IC) as a very weak implicature, perhaps as something on the borderline between implicatures and completely unintended contextual implications of the utterance. This is due to the fact that non-target consumers are those who did not learn to associate "malore improvviso" to a deadly outcome of Covid-19 vaccination. So in the comprehension procedure of (8), (8A) and (8B) are useless for them, because it cannot produce any context expansion.

Secondly, the inference of (8IC) gives the target group the possibility to recover an additional range of progressively weaker implicatures:

- If Caterina had not been vaccinated, she would probably still be alive.
- Caterina had not been warned about the high risks of vaccination.
- Caterina was forced to get vaccinated.

The author of (8) did not provide enough evidence for the recovery of such implicatures: there is hardly anything in (8) which may suggest that Caterina was forced to get vaccinated, for instance (also because the headline does not even claim that Caterina got vaccinated against Covid-19). However, recovery of further implicatures depends also on the kind of encyclopaedic information the hearer has access to. If someone has learned to associate the information that a sudden illness is a probable

outcome of Covid vaccination, then they may have access to this chunk of information, through context expansion, every time the expression "sudden illness" (in our case "malore improvviso") pops up. Again, accessibility is not always guaranteed: other conditions might be satisfied too. In our example, I hypothesized that the context of interpretation of (8) included at least the assumptions (8A) and (8B), which are useless and inoperative for non-target consumers, but are instead fruitful for target consumers[74].

Why is news with covert target fake news? Contrary to what one may think reading the analysis of (8), the problem is not really the fact that headlines like that may induce the recovery of false implicatures in the target group. I regard this (despicable) fact as an instrumental way to exploit news with covert target, but I do not think that this is the reason why that kind of news is fake. In fact, regarding deceptive headlines, I would still classify news with cover target as fake news also when it furtively conveys true implicatures. The wrong in targeted fake news is that it violates both norms of production and norms of publication.

Firstly, when the target of news is covert there is a portion of the audience, the non-target consumers, which is not completely enabled to keep track of what is reported. By definition, the non-target group cannot have access to the same inferential path as the target group, which benefits of different, and sometimes richer, cognitive effects. It is as if the same piece of news reaches the second group in its entirety, but not the first one, which is in a sense excluded from the information exchange.

But this sets news with cover target apart from real news conceptually. In fact, while the first tends to exclude certain members of the audience from the information exchange, the second is always accessible to all readers – even when the content is not interesting for any news consumer, as in the case of local news. In fact, one essential duty of journalism, the social institution that has a legitimate monopoly on news production, is serving the interest of citizens in general (Kovach & Rosenstiel 2014). Ideally, a journalist has no interest in reaching only a partition of the audience and excluding the rest: real news producers have no covert target groups because they

[74] The function of target fake news is akin to the one of so-called "dogwhistles" in philosophy of language (Witten 2014, Saul 2018, Henderson and McCready 2019, Mascitti 2023). Both dogwhistles and targeted fake news work via exploitation of diverse beliefs among various subsets of the audience (Khoo 2017), with the goal of getting across different messages with the same utterance and at the same time. However, in the next chapter we will see a case of targeted fake news that, contrary to dogwhistles and the headline (8) here, is not driven by the use of a specific expression.

(should) have no agenda other than commitment to the truth. By excluding a portion of the audience from the information exchange, news with covert target is violating a fundamental norm of production: the commitment to the interest of *all* citizens.

Secondly, producers of news with covert target do not respect norms of news publication, because they do not take full responsibility for what they communicate in a communicative context where there is no mutual adjustment. As we have already mentioned, when the exchange of information occurs in the absence of a context of mutual adjustment, the communicator bears more responsibility than they would in a normal context. By addressing a covert target and playing on the encyclopedic information associated with the concepts encoded by certain expressions, the producer is not taking full responsibility for what the reader will understand, because additional implicatures like (8IC) and the weaker ones are, at most, encouraged but not intended. Formally, in fact, the producer does not provide any evidence of intending to communicate (8IC), and thus does not take responsibility for it.

But if the producer is actually aware that headlines like (8) will have this effect on some portion of the audience – the covert target – and does nothing to prevent it, then they are failing to meet the responsibility required in that context, where they know the news consumer will not be able to ask for clarification if uncertain about the understanding (e.g., they cannot ask to explicitly cancel (8IC) or the others)[75].

Conclusion

In this chapter, I put forward a new definition of fake news, according to which fake news is news that does not deserve its status. To unpack all the implications of this definition, I proposed two separated, but equivalent, interpretations of fake news: FN_A and FN_B. They are equivalent because the expectations of news consumers around news is conceptually related to the observance of journalistic norms: news that violates norms of journalism does fall short of expectations around a journalistic product.

This new definition of fake news is intended to be part of an institutional account of fake news which interprets the element of fakeness in fake news as the attempt of some

[75] Guessing intentions is always hard work, but if we consider the continuous omission of a rectification amid the throng of misguided claims in the comment section, we can reasonably conclude that the authors of headlines in the style of (8) were probably aware of the likely effects of their utterances.

source of information to boast about having the status of a journalistic source, when actually they do not have it. The institutional account of fake news conceives fake news as conceptually separated from pathologies of information, thus avoiding the objection of conceptual necessity from Habgood-Coote (2019). “Fake news” does not denote a kind of flawed information – like false or misleading information – but refers to news that does not comply with the norms of journalism. Accordingly, news that is false is not necessarily fake news (and vice versa), and boundary categories like “false news”, “misinformation”, and “disinformation” refer to distinct phenomena that can only potentially overlap with those denoted by “fake news”. For instance, if a piece of news that significantly violates norms of journalism encapsulates also false information, then we may say that this is a case of fake news *and* misinformation.

Finally, the introduction of the norms of news publication has allowed us to identify new phenomena of fake news that would not emerge from other accounts and to provide an alternative interpretation of clickbaiting (no longer tied to the concept of bullshitting). Using the theoretical tools of Relevance Theory, I analyzed three classes of headlines: clickbaiting headlines, deceptive headlines, and targeted headlines (or targeted fake news). The analysis of these particular types of fake news will be useful for applying the institutional account to a specific news topic: climate change.

The next and final chapter of my thesis aims to show how the use of the institutional account of fake news helps us capture new phenomena in the specific field of climate change.

Chapter 4

Fake news about climate change: what's new?

1. What is climate change?

Generally speaking, climate change refers to long-term alterations in temperature, precipitation, wind patterns, and other elements of the Earth's weather system, primarily (but not solely) characterized by a significant increase in global average temperatures. This increase is largely driven by the accumulation of greenhouse gases – like carbon dioxide, methane, nitrous oxide, and water vapor – in the atmosphere, resulting from human activities such as the combustion of fossil fuels, deforestation, and industrial processes (Rogelj et al., 2016). The phenomenon has far-reaching implications for natural and human systems, affecting ecosystems, agriculture, water resources, and public health.

Rising temperatures contribute to more frequent and severe weather events, such as hurricanes, droughts, and floods (Coumou & Rahmstorf, 2012). These phenomena not only threaten human life and infrastructure but also disrupt ecosystems, leading to loss of biodiversity (Doney et al., 2012). Ocean temperatures are also increasing, causing coral bleaching and altering marine habitats, which has dire implications for global fisheries and food security (Hoegh-Guldberg et al., 2018). Moreover, climate change exacerbates social and economic inequalities. Vulnerable populations, particularly in developing countries, are often the most affected due to their limited resources and adaptive capacities (Hallegatte et al., 2016). Displacement due to rising sea levels and extreme weather events can lead to climate refugees, intensifying geopolitical tensions and straining national and international resources (Ferragina & Quagliarotti 2019).

Apart from its everyday use, there are various competing ways to interpret the concept of climate change. As Frigg, Thompson and Werndl (2015) maintain, the reason is that climate itself is an elusive concept. Intuitively, the climate is the distribution of certain variables (called climate variables) arising for a certain configuration (i.e. certain gas greenhouse concentrations and certain aerosol emissions) of the climate system. The climate variables include those that describe the state of the atmosphere and the ocean and sometimes also other variables such as those describing the state of glaciers and ice sheets. So the climate is not about the exact values of the surface air temperature, ocean

temperature etc. at a certain point of time, but about the surface air temperature, ocean temperature etc. that one can expect when the climate system is in a certain configuration (Frigg et al. 2015).

But Frigg and his colleagues notice that there are two different kinds of definitions of climate, climate as "distribution over time" and climate as "ensemble distribution":

> Climate as *distribution over time* roughly corresponds to what one learns about climate in geography education at school. Here, climate is about the distribution of the climate variables that arises when the climate system evolves over a certain time period. Climate as *ensemble distribution* is popular with scientists who are concerned with predicting the climate. It is about the distribution quantifying how likely it is that certain values of the climate variables are expected at a certain point of time in the future given our current uncertainty about the climate variables. (Frigg et al. 2015: 953)

There are several definitions of both climate as distribution over time and climate as ensemble distribution[76], and each of them suffers from certain limitations. The former group of definitions usually fail to reflect actual climate behavior, because they either underestimate the fluctuations of real-world external conditions (like volcanic activity and solar radiation) or do not account for regime shifts, that is sudden and profound changes in the climate system, such as an ice age or a meteor impact. The latter group of definitions might be better equipped for reflecting climate behavior, but nonetheless are either too distant from our intuitions of what climate is or fail to address how we should define past or present climates.

To avoid complications, and since providing the ultimate definition of climate and climate change goes beyond the scope of this thesis, I am going to stick with the definitions on the same matters adopted by the Intergovernmental Panel on Climate Change (IPCC):

> CLIMATE: Climate in a narrow sense is usually defined as the average weather, or more rigorously, as the statistical description in terms of the mean and variability of relevant quantities over a time period ranging from months to thousands or millions of years. The classical period for averaging these variables is 30 years, as defined by the World

[76] For definitions of climate as distribution over time, see Lorenz 1995; Dymnikov and Gritsoun 2001; Werndl 2015. For definitions of climate as ensemble distribution, see Stone and Knutti 2010; Daron and Stainforth 2013.

> Meteorological Organization. […] Climate in a wider sense is the state, including a statistical description, of the climate system. (IPCC: 2013: 126)
>
> CLIMATE CHANGE: A change in the state of the climate that can be identified (e.g., by using statistical tests) by changes in the mean and/or the variability of its properties and that persists for an extended period, typically decades or longer. Climate change may be due to natural internal processes or external forcings such as modulations of the solar cycles, volcanic eruptions and persistent anthropogenic changes in the composition of the atmosphere or in land use. (IPCC 2019: glossary)[77]

IPCC's definitions are also the gold standard for most of the experimental studies on fake news and mis/disinformation on climate change, which are the focus of this chapter. Here is how it proceeds.

In the next section (§2), I will summarize the main findings about disinformation on climate change (or climate disinformation), highlighting its role in fostering scepticism towards climate science. Then, in §3, we wade into the question of whether there is a problem with climate fake news too, and we are going to investigate if such notion is sufficiently disentangled from climate disinformation or not. The core of this chapter, though, is the section §4. Here I apply the institutional account developed in the previous chapter to the case of climate fake news, checking whether we can fruitfully use this approach to diagnose more interesting cases of fake news or rule out less controversial ones.

2. Disinformation: the main enemy of climate science

Climate change is widely accepted within the scientific community as a critical and urgent issue. Since the 19th century, when the physicist John Tyndall demonstrated that CO_2 absorbs heat, the scientific understanding of climate change has steadily progressed. The evidence accumulated over decades now unequivocally shows that human activities, as we said in §1, are the first driver of the increase of greenhouse gases in the atmosphere, driving global warming and other associated impacts that we subsume under the general label of "climate change" (Mitchell et al. 2001; Allen et al.

[77] The definition of climate change given by the IPCC belongs to the second group of definitions listed before, those that interpret climate as an ensemble distribution.

2006; North 2013; AR4 2007; AR5 2014; AR6 2021). Therefore, it now became a platitude among journalists, pundits, and independent commentators to use "climate change" as an elliptical expression in the place of "anthropogenic climate change".

However, despite the strong scientific consensus – where nearly 97% or more of climate scientists agree on human-caused climate change (Cook et al. 2013, 2016; Powell 2019; Lynas et al. 2021) – public perception, especially in countries like the United States, has not fully aligned with the received view of climate scientists. This misalignment is not simply due to a lack of information but is also driven by political polarization and well-orchestrated disinformation campaigns (Lewandowsky 2021).

In fact, one of the major challenges in public acceptance of climate science and subsequent acts of mitigation is that all these things often conflict with deeply held worldviews, particularly among those with libertarian or conservative ideologies who perceive climate mitigation as a threat to free markets (Lewandowsky et al. 2013; Hornsey et al. 2016; Keith Smith and Mayer 2018). Addressing climate change requires policies that regulate emissions, tax carbon, or limit the use of fossil fuels, all of which challenge laissez-faire economic principles. As a result, climate scepticism is strongly associated with right-wing political ideologies, particularly in countries like the United States, Australia and the United Kingdom. The situation in continental Europe is no different according to recent data (Kulin et al. 2021).

This political polarization of climate change has created a fertile ground for the rise of climate disinformation, where vested interests exploit ideological divides to prevent meaningful action. Disinformation about climate change, in fact, is not a spontaneous or grassroots phenomenon. Research shows that it is actively and systematically disseminated by well-funded networks (Lewandowsky 2021). These networks have been traced to conservative think tanks, many of which receive substantial funding from fossil fuel companies like ExxonMobil and organizations tied to the Koch brothers (Farrell 2016b). The financial backing of these actors allows them to flood the media with disinformation, much of which is aimed at casting doubt on climate science[78].

[78] Those networks may include several different actors that do not boil down to think tanks alone. They can involve isolated group of scientists serving the interests of specific companies – like the handful of ill-intentioned physicists who are present in Oreskes and Conway (2010); governments; foundations; industrial companies, usually specialized in oil or coal extraction; media, particularly those with right-wing affiliations. For more details, see Bjornberg et al. (2017).

For example, as Lewandowsky (2021: 6) reports, more than 90% of environmentally sceptical books in the United States can be linked to conservative think tanks, and most of these books have not undergone any peer review. Between 2003 and 2010, these think tanks received approximately $900 million annually, although it is unclear how much of this money is specifically used for climate disinformation. Additionally, lobbying efforts by these industries have been substantial, with over $2 billion spent on lobbying Congress from 2000 to 2016 to block climate legislation. These financial resources have made climate disinformation a formidable barrier to public understanding and political action.

The climate disinformation strategy is particularly effective also because of its influence on mainstream media. The longstanding journalistic practice of presenting 'both sides' of an issue sometimes has been manipulated to create a false equivalence between climate scientists and sceptics (Boykoff 2008). This inappropriate balance gives disproportionate coverage to the few contrarian voices, leading to public confusion about the state of art on climate change within the scientific community. It has been demostrated that inappropriate balance in scientific discussions reduces public's understanding of both the amount of available evidence and the scientific consensus around the same issue (Koehler 2016, Cook et al. 2017).

Although the media's coverage has improved in recent years (Brüggemann and Engesser 2017), the legacy of inappropriate balance continues to affect public understanding. Moreover, conservative media outlets, such as Fox News in the United States, continue to disproportionately amplify climate denial narratives (Bolin and Hamilton 2018; Gustafson et al. 2019). Social media platforms have further exacerbated the problem by enabling the rapid dissemination of climate disinformation. For example, most climate-related YouTube videos promote antiscientific views, and the most widely shared article on climate change on social media in 2016 was a piece of disinformation that falsely claimed that 31,000 scientists disputed the consensus on climate change (Readfearn 2016).

For all this reasons, in the last decade the scientific literature addressing the issue of climate scepticism has been focusing prevalently on the phenomenon of disinformation. Concerning climate change, it makes more sense to talk about disinformation rather than misinformation or other pathologies of information, for we said that false

information on the matter was spread in an organised fashion. Most of the time the attempts to discredit climate science were systematic and premeditated; they were not the by-products of malpractice or inexperienced online users.

The literature on climate disinformation has also produced an array of different terms to describe the organized campaigns brought about think tanks, industries and hyperpartisan media. Treen and colleagues (2020) report many examples of this terminology: "sceptical discourse" (Boussalis and Coan 2018), "sceptical arguments" and "denial campaign" (Elsasser and Dunlap 2013), "attempted expert knowledge de-legitimization and contestation" and "reinterpretation of existing climate science knowledge claims" (Sharman 2014), "contrarian messages/discourse" (Farrell, 2016a), "contrarian information" (Farrell, 2016b), and "scepticism and doubt" (Reed 2016). This terminology points towards the same phenomenon: systematic attempts by various actors to discredit climate science and confuse political debate on climate change, by casting doubt on well-supported theories and providing alternative and often fallacious interpretations of empirical observations.

The main rhetorical strategies of these actors are almost always the same. Concerning disinformation that attempts to simply deny the findings of climate science, Bjornberg et al. (2017) have identified four widespread strategies: the trend denial, which consists in claiming that no significant warming is taking place on our planet; attribution denial, that rejects the anthropogenic nature of climate change in order to remove responsibility from national governments and polluting companies; impact denial, that minimizes the significant impact of climate change on humans or the environment; and consensus denial, that attempts to obscure the vast consensus among climate scientists about anthropogenic climate change.

But simple denial is not the only strategy adopted by actors of climate disinformation. Other strategies include stressing uncertainties and demand absolute certainty as a necessary condition for climate action, attacking scientists to undermine their credibility (for instance, through the imaginary reconstruction of events and facts around the so called "Climategate scandal"), delegitimizing scientific practices and institutions in general (such as the peer review), and issuing pseudoscientific alternatives to climate science (Lewandowsky 2021: 7).

One of the most damaging aspects of climate disinformation is its ability to generate confusion over already settled debates. Even a handful of misleading statistics can override the acceptance of accurate information about global warming (Ranney and Clark 2016). Furthermore, when people are presented with two opposing viewpoints, as in the case of inappropriate balance debates, they often view them as equally valid, leading to climate inaction. A particularly notable example is the widely shared petition that falsely claimed 31,000 scientists rejected climate consensus (already mentioned above, p. 182); this document, despite its dubious credibility, was a major blow to public acceptance of climate change.

Besides, climate disinformation has a profound impact not just on public opinion but also on political attitudes and scientific communication. The barrage of contrarian activities targeting individual scientists, often through public criticism, affects how researchers communicate and can delay knowledge dissemination. Scientists, pressured by external attacks, may adopt cautious language, as seen in the IPCC reports, which tend to express moderate levels of confidence even when evidence is overwhelming (Medimorec and Pennycook 2015).

3. Is there a problem of climate fake news?

As we said so far, the public discussion on climate change has been mudded by organised disinformation campaigns. Many scientists believe that disinformation is still the main enemy of climate science and a great obstacle for the public acceptance of policies that may slow down or even arrest, in the long run, the trend of global warming. But what about fake news? Can we say that there is also a problem of fake news around the topic of climate change (henceforth, climate fake news)? And if so, how do we assess it? What kind of phenomena does it capture that a simple notion of disinformation cannot?

Answering these questions is not an easy task, for several reasons. First of all, if the philosophical community – together with the scientific one – does not agree on a common definition of fake news, it is much harder to establish whether we have a problem with fake news or not, regardless of the topic under discussion. The reason is that, as long as each scholar (or groups of scholars) relies on a different definition of the same notion, they will look for different pieces of evidence. For instance, the absence of

a common definition of fake news may impinge upon the collection of datasets for empirical studies that have the same research question, because different teams of scientists may include discordant things under the same label of “fake news”. At that point, it becomes difficult to compare the results of such studies.

Luckily, studies on climate fake news exhibit a certain level of coherence in the definition adopted. As we shall see soon, the scientific literature about climate fake news is scant. However, scholars who investigates the detrimental role of fake news within climate change discussion adopt prevalently one of two privative definitions. The most widespread is the definition offered by Lazer et al. (2018), which states that fake news is fabricated information that mimics news media content in form but not in organizational process or intent. This definition is privative and highlights the intentional fabrication of information that appears like legitimate news but lacks the standards of journalistic processes. However, while Lazer’s framework is widely used, the field is nuanced, and various studies adopt or build on this definition to address the complexity of fake news, especially concerning climate change.

The other adopted definition comes from the work of Tandoc et al. (2018). Tandoc, Lim, and Ling categorize fake news into different types based on intent and connection to factual content, which helps distinguish fake news from other misleading information forms. By making a review of 34 different articles, they highlight that what is common across each article is that «fake news appropriates the look and feel of real news» and «hides under a veneer of legitimacy as it takes on some form of credibility by trying to appear like real news» (Tandoc et al, 2018: 11).

An important aspect of this investigation is the disparity between studies on climate disinformation and those on climate fake news, with the former being far more prevalent. To explore this, I conducted a Scopus search focusing on disinformation and fake news related to climate change. For climate disinformation, I queried “misinformation” or “disinformation” paired with “climate change” or “global warming” and excluded “fake news”[79]. This query accounted for interchangeable usage

[79] Here you find the advanced query: (TITLE-ABS-KEY ({misinformation} AND {climate change}) OR TITLE-ABS-KEY ({disinformation} AND {climate change}) OR TITLE-ABS-KEY ({misinformation} AND {global warming}) OR TITLE-ABS-KEY ({disinformation} AND {global warming}) AND NOT TITLE-ABS-KEY ({fake news})) AND (EXCLUDE (SUBJAREA , "DENT") OR EXCLUDE (SUBJAREA , "IMMU") OR EXCLUDE (SUBJAREA , "PHYS")) AND (EXCLUDE (DOCTYPE , "er")). Someone could be tempted to arbitrarily select those subject areas that

of “misinformation” and “disinformation” and older studies referencing “global warming.” It yielded 463 results, later refined to 457 after excluding minor subject areas (e.g., dentistry, physics) and document types like errata. The documents span from 1998 to 2024. For climate fake news, I queried “fake news” with “climate change” or “global warming,” excluding “disinformation” and “misinformation”[80]. This search returned 74 results, refined to 55 after filtering, with studies dating back to 2017—post-2016, when the term “fake news” gained prominence. The final count revealed 457 documents on climate dis/misinformation versus 55 on climate fake news.

While not conclusive evidence of a lack of interest, this highlights a significant focus on disinformation over fake news in the scientific community.

3.1 Climate fake news: what we know so far

We said that most of the studies on climate fake news adopts a privative approach: fake news (on climate change) is fabricated information that resembles real news in appearance but not in the respect of journalistic standards. But what are the main findings of this literature? What do these studies present as the main characteristics of climate fake news? I will consider a handful of studies whose findings are sufficiently representative of this literature, privileging the most cited ones.

It must be clear by now that we are not interested in considering scientific literature that conflates climate fake news with climate dis/misinformation, but we only review that tiny subset of works which acknowledges an independent existence to the phenomenon of (climate) fake news. Therefore, it makes no sense to report here studies that use “climate fake news” in a non-technical sense, as a synonim for climate disinformation. Most of the time these studies end up joining the wider literature on

---

seems to be intuitively distant from the fields that study climate disinformation and filter them out (like “nursing” or “agricultural and biological sciences”). But more often then not, you find relevant studies about the same issue that are indexed as part and parcel of an unexpected subject area. To avoid this problem, I opted for a more inclusive query.

[80] (TITLE-ABS-KEY ({fake news} AND {climate change}) OR TITLE-ABS-KEY ({fake news} AND {global warming}) AND NOT TITLE-ABS-KEY ({misinformation}) AND NOT TITLE-ABS-KEY ({disinformation})) AND (EXCLUDE (SUBJAREA , "ENGI") OR EXCLUDE (SUBJAREA , "PHYS") OR EXCLUDE (SUBJAREA , "DECI") OR EXCLUDE (SUBJAREA , "BIOC") OR EXCLUDE (SUBJAREA , "AGRI") OR EXCLUDE (SUBJAREA , "PSYC") OR EXCLUDE (SUBJAREA , "NEUR") OR EXCLUDE (SUBJAREA , "ENER") OR EXCLUDE (SUBJAREA , "CHEM") OR EXCLUDE (SUBJAREA , "NURS") OR EXCLUDE (SUBJAREA , "MULT") OR EXCLUDE (SUBJAREA , "MATE") OR EXCLUDE (SUBJAREA , "EART")).

climate disinformation that we have summarized before, so they fall outside the current objects of this section.

One prominent line of research on climate fake news is its impact on beliefs about climate change, scientific consensus, and trust in scientists. In an experiment conducted by Drummond et al. (2020), participants were exposed to false news headlines casting doubt on anthropogenic climate change under several conditions:

1. Control: unrelated fake news headlines.
2. Warning: a pre-exposure warning about the falsity of the climate change-related fake news.
3. Debunking: a post-exposure statement that the headlines were false.
4. Uncorrected: exposure to fake news with no warning or debunking.

Following the privative approach of Lazer et al. (2018), the authors of the experiment considered only information that mimics credible reporting in format but not in content or intent to be fake news. They found that exposure to fake news did not significantly alter overall beliefs in climate change across different groups. Those who received the fake climate news (without warnings or corrections) reported slightly lower beliefs in anthropogenic climate change than the control, but the effect size was minimal.

At the same time, there was a small reduction in participants' estimates of the scientific consensus on climate change when exposed to fake news. The warning or debunking interventions did not fully mitigate this impact, with some evidence suggesting that the warning condition might actually have reinforced scepticism.

Trust in scientists remained relatively unaffected across conditions. However, in the debunking condition, participants exhibited marginally lower trust in scientists compared to the control, a phenomenon that scholars call "backfire effect": attempts to correct people's misguided belief actually results in the reinforcement of those same beliefs (Nyhan and Reifler 2010; but see Wood and Porter 2019 for a contrarian view).

The authors also found that political ideology and affiliation were more significant predictors of climate change scepticism and trust in scientists than fake news exposure. Republican and conservative participants reported less belief in climate change and lower consensus estimates compared to Democrat and liberal participants, reinforcing the polarization of climate beliefs along ideological lines.

In conclusion, Drummond and colleagues claim that exposure to climate-related fake news has limited effects on shaping climate scepticism or altering beliefs in scientific consensus. Political ideology is a stronger determinant of climate change beliefs than fake news exposure. Given the small effect sizes observed, the authors suggest that fake news might have more of a symbolic or identity-signaling role rather than a direct influence on beliefs.

The tight relation between climate fake news and political ideology has been highlighted by other studies and represents another prominent line of research on the topic. For instance, Al-Rawi et al. (2021) analyzed 12.055 tweets from a larger set of 6.8 million posts referencing "fake news", focusing on those that also mentioned climate change or global warming[81]. Al-Rawy and colleagues did not adopt any specific definition of "fake news", since their idea is that, when it comes to media discussions of climate change, the term fake news ceases to be a signifier and merely becomes a tool to attack opponents.

To prove this, their study categorized tweets based on belief in climate change (anthropogenic vs. denial) and political tone (anti-liberal or anti-conservative). Tweets labeled as anti-liberal or anti-Democratic appeared twice as often as anti-conservative or anti-Republican ones (24.2% vs 12.4%). Much of the discourse leveraged climate change scepticism to critique left-leaning policies and figures, illustrating the politicization of climate science. For instance, popular tweets disparaged media outlets like CNN and liberal figures, including Congresswoman Alexandria Ocasio-Cortez, linking them to perceived misinformation about climate impacts like the Australian bushfires.

Additionally, the study identified Donald Trump as the most mentioned Twitter user, with many posts either supporting or echoing his frequent use of "fake news" to dismiss media reports critical of his stance on climate issues. Additionally, media outlets, especially CNN, were frequent targets of these tweets, labeled as promoting "fake news" regarding climate change.

The study of Al-Rawi et al. thus concludes that the term "fake news" is a hollow one that has become a politicized tool within climate change discussions, primarily used by

[81] Notwithstanding the fact that Twitter formally does not exist anymore, I still talk about "tweet/s" here because Al-Rawy and colleagues had conducted their study before Elon Musk took ownership of the platform (now called "X").

right-wing and anti-liberal Twitter communities to attack mainstream media and climate science. Also, since the authors found that among the top 500 retweeted posts 55,8% denied climate change or attributed it to natural cycles rather than human activity, they concluded that dominance of climate denial (or trivialization) in top tweets suggests that social media amplifies polarized views, allowing ideologically-driven actors to shape public discourse around environmental issues.

The role of political ideology in the acceptance of climate fake news is further stressed by other studies. Lutzke et al. (2019) focused on testing interventions that encourage critical thinking as a method to reduce belief in and sharing of fake news about climate change on Facebook. The definition of fake news adopted by Lutzke and colleagues comes again from Lazer et al. (2018), but here it is slightly altered: fake news is not only fabricated information which mimics the appearance of legitimate reporting, but it is also intended to mislead. Even though the authors make an explicit reference to the work of Lazer and colleagues, their adopted definition ends up being a mixture of privative and conservative accounts.

The study aimed to examine whether two short, simple critical thinking prompts could help Facebook users better distinguish between fake and real news on climate change. Here are the two kinds of prompts:

1. Guidelines Condition: Participants read a brief set of guidelines on how to evaluate online news, focusing on source credibility and checking for political motivation.
2. Enhanced Guidelines Condition: Participants not only read the guidelines but were also asked to rate the importance of each guideline, promoting deeper reflection.

Nearly 2750 participants were then divided among guidelines, enhanced guidelines and a control condition. Each group was exposed to fake or real Facebook news posts about climate change, taken from either hyper-partisan sources (like *InfoWars* and *Breitbart*) or reputable sources (like *USA Today*). They rated the credibility of each post and their likelihood of liking or sharing it, simulating typical Facebook engagement.

Both prompts were effective in reducing trust, liking, and sharing of fake news posts. However, the effect size was modest, suggesting that while prompts can be effective,

their impact is limited on social media[82]. But once again, conservative participants showed greater trust in and likelihood of sharing fake news posts compared to liberal participants, a trend associated with motivated reasoning[83]. Conservatives in the study tended to be more sceptical of anthropogenic climate change, often perceiving climate-related information from hyper-partisan sources as aligning with their views.

Even among conservative participants, the interventions had a measurable effect in reducing trust in fake news, though the effect size was smaller than for more liberal participants. This suggests that while critical thinking prompts can reduce engagement with fake news across ideological divides, the effect is moderated by participants' political beliefs. Notably, conservatives with higher climate knowledge were less susceptible to fake news, showing that knowledge can offset some ideological biases.

So what do we know about climate fake news so far? What immediately leaps out from this brief overview of some of the most cited studies on the issue is the strong similarity with findings around climate disinformation. The acceptance and even the propagation through re-sharing of climate fake news and climate disinformation is strongly influenced by political ideology. In both cases, a libertarian/conservative ideology is linked to a higher propension to believe in the content of climate fake news and climate disinformation. In other words, both phenomena primarily impact those who are already sceptical of climate science.

There may be subtle differences concerning the authorship. Disinformation is typically crafted intentionally by actors with specific goals, like fossil fuel interests or political groups aiming to protect the status quo. Fake news, in contrast, can originate from sources seeking to gain social media engagement, ad revenue, or attention without necessarily pursuing a broader ideological agenda. The study by Al-Rawi et al. (2021), presented before, supports this view on fake news. But it is also true that the authors of this study did not commit themselves to a specific definition of fake news, leading to a broader inclusion of cases. If they had chosen a privative definition, like the one by Lazer et al. (2018), they would have been forced to leave out tweets that did not presented information with the façade of legitimate media.

[82] Furthermore, Lutzke and colleagues found that both conditions with guidelines did not reduce at the same time trust and propension to share real news on climate change. This is important because it means that use of guidelines has targeted fake news selectively.

[83] Motivated reasoning refers to biased information processing in accordance with prevailing motivations and worldviews (Kunda 1987).

Furthermore, the studies presented so far show that both climate disinformation and fake news manipulate cognitive biases, appeal to emotions, and exploit media dynamics to influence public perception and delay policy support. Their similar impacts reflect their shared capacity to reshape climate discourse by undermining trust in science, amplifying ideological divides, and promoting misinformed public and policy perspectives. This overlap in psychological appeal and societal impact is why scientists often study them in similar ways, even as they acknowledge distinct differences in intent and origin.

However, as we well know from the previous chapters, disinformation and misinformation are more illuminating and helpful terms than "fake news". Disinformation refers to false or misleading information deliberately created to deceive, while misinformation refers to the spread of false or misleading information without deceiving intent. Fake news, on the other hand, has become a broad, somewhat ambiguous term, that in many contexts has been also weaponized. But scientists favour precise language that clearly defines whether the intent behind a piece of information was malicious or accidental and prefer to avoid politically laden concepts.

For instance, here is how Soroush Vosoughi, Deb Roy, and Sinan Aral commented on the issue of terminology:

> Although, at one time, it may have been appropriate to think of fake news as referring to the veracity of a news story, we now believe that this phrase has been irredeemably polarized in our current political and media climate. As politicians have implemented a political strategy of labeling news sources that do not support their positions as unreliable or fake news, whereas sources that support their positions are labeled reliable or not fake, the term has lost all connection to the actual veracity of the information presented, rendering it meaningless for use in academic classification. We have therefore explicitly avoided the term fake news throughout this paper and instead use the more objectively verifiable terms "true" or "false" news. (Vosoughi et al. 2018: 1146).

Since in scholarly work impartiality and specificity are crucial, and seeing as how "fake news" has acquired over time politically charged and subjective connotations, often used as a pejorative by various political groups, scientists have an additional reason to opt for other, less polarized terms. Terms like "false news" or "disinformation" allow

scientists to objectively analyze the phenomenon without the term itself implying a specific stance or political angle.

These observations on meaning instability and political weaponization, combined with a substantial overlap in central findings, may be some of the possible causes of the disproportion between the amount of research on climate disinformation and the one on climate fake news. If so, it might be worthwhile to use an alternative definition of fake news that can yield more interesting results in the research on climate fake news.

4. **Institutional climate fake news**

In this section we deploy the institutional account of fake news, developed in the previous chapter, to investigate the phenomenon of climate fake news. The goal here is to show how this alternative account may either pin down the real wrong in some, already considered, fake news or identify new cases of fake news that the existing accounts cannot properly account for.

The structure of this section reflects that of §4 in the previous chapter. We are going to present different kinds of violation of the norms of journalism by news dealing with the issue of climate change. We will consider violations that emerge from the reading of headlines or from a combination of the headline and the subsequent text of the article. In the third chapter, here you find an investigation of three class of problematic headlines: clickbait headlines, deceptive headlines and targeted headlines (or targeted fake news).

When necessary, the context of the news will be reconstructed, along with an explanation of why certain informational content is false. However, it is important to remember that what is problematic about fake news – as interpreted within an institutional account – is not their (alleged) falsity, but rather the degree to which they violate the norms that confer the status of news upon information. This means that a false story on the topic of climate change is not necessarily an example of (institutional) climate fake news; but it may sometimes count as climate disinformation or misinformation, at most.

4.*1 Clickbait headlines on climate change*

Although there is extensive literature on clickbait headlines and how they function, few analyses focus specifically on clickbait headlines on climate change. One of them is the study of Xu et al. (2022), which analyses the role of emotions in clickbait headlines on the topic of climate change.

According to their study, clickbaiting in climate change headlines uses a set of emotional, linguistic, and structural strategies aimed at maximizing reader engagement and shareability. Primarily, these headlines leverage emotional appeals, particularly high-arousal emotions such as fear, anger, and surprise. Xu and colleagues suggest that when climate change headlines evoke fear – through terms implying urgency or catastrophe – they attract more reader attention, especially for articles that *support* climate change science[84].

Moreover, the research highlights that negative framing in headlines – such as using terms like "threat," "disaster,"– is especially common in denial articles and increases engagement by amplifying scepticism and doubt. Denial headlines also frequently use language tied to anger or disbelief, playing into partisan divides and reinforcing echo chambers. The study references previous work showing that negative and sensational language in clickbait tends to outperform neutral or positive phrasing, a tactic that denial content often adopts more aggressively than pro-science articles.

However, the study of Xu and colleagues sometimes seems to conflate emotionally laden headlines with clickbait headlines, because both are geared to the maximization of the engagement. But in the previous chapter we said that a defining features of clickbait headlines is their forward-referencing nature: they build a bridge between them and the article, so that consumption of the latter becomes a necessary requirement for the cognitive processing of the former. So the appeal to emotions is not per se sufficient to have a clickbait headline; there must also be the tendency of the headline to point to its accompanying text.

For instance, Xu and colleagues claim that the structure of climate clickbait headlines often employs short, punchy formats, occasionally phrased as provocative questions or

[84] This emotional pull operates through what's known as a "dual-systems model" of cognition (Cacioppo et al. 1999; Thayer 1989; Watson et al. 1999), where instinctive emotional responses prompt rapid, sometimes subconscious, engagement decisions, driving clicks before readers engage critically with content.

exaggerated claims, which creates a sense of curiosity or urgency that tempts readers to click. Additionally, ambiguous or suggestive phrasing can entice readers by hinting at dramatic or hidden insights without fully disclosing the content, a technique that plays well with high-arousal emotional appeals. But even though emotions may play a role in arising curiosity over a topic, this is not the same as forward-referencing, because emotionally laden headlines are still autonomous texts: readers do not need to access the accompanying the text to process them cognitively. To understand the difference, consider this pair of headlines:

9) White House: Climate Change Is Real And It's Making Your Life Wors Right Now. (McMorris-Santoro 2014)

10) Climate Change Catastrophes Ahead? Here's How We Can Prepare. (*NBC News*)

Both headlines follow the release of a White House's report on climate change in the United States. The report was issued in May 2014 and was covered by many news media outlets with different presentations. A steady reaction of media was one of fear: the summaries of the report were focusing over its terror-inducing passages about the major effects of climate change over the lands, the economy and the overall future of United States. The headlines reporting the news were reflecting this communicative strategy, but not all of them were clickbait.

(9) is clearly an emotionally laden headline about the report. It evokes a potent sense of fear and addresses the reader directly ("your life") as if it was an admonishment. We could discuss whether (9) violates norms of journalism to a sufficient extent; surely it does not represent an impeccable example of journalism. But (9) is not a clickbait headline, because it is not forward-referencing. The cognitive effects of (9) are sufficiently achieved through the reading of the headline alone; you do not need to read the accompanying text too.

It is totally different for (10), which is instead a clickbait headline for all intents and purposes. Like (9), (10) adopts a language that promotes fear around the topic of climate change ("Climate Change Catastrophes"). But contrary to (9), the reading of (10) leaves a sense of cognitive 'dissatisfaction' that is typical of clickbait. The use of "Here" invites the reader to click through the link and go to the landing site, because you cannot fully grasp the informative contribution of (10) without reading the

accompanying text. In fact, in the previous chapter (§4.2) we said that demonstratives are used to point something or someone out to the reader. But demonstratives are not the only class of terms that appears in clickbait headlines. There is also a wide array of indexical expressions that may serve the same crucial role; "Here" is one of them.

Generally, "here" stands for the place where this word is pronounced, but in an online title it refers to the place where the text accompanying this title is located. So using "here" in the headline means pointing toward a specific piece of information that the reader reasonably expects to find only through the reading of the accompanying article. In doing so, the author of (10) explicitly directs the reader's focus to the strategy by which we can face the climate crisis that lies ahead in the future; a strategy that remains unknown in the headline. To use the lexicon of RT, the use of adverb "here" raises our expectations that following the link will be optimally relevant, because we have a chance to know what the strategy is. We may also speculate that the negative emotions prompted using "catastrophe" amplify the cognitive need to discover the strategy to face such catastrophe, but the extent of this effect, if any, remains unknown.

It is interesting to note how, in the case of climate change, clickbait headlines are more commonly used by accredited news outlets rather than by blogs associated with disinformation on the topic. For example, Fuller (2014) provides a list of clickbait headlines published following the White House report in May 2014, which we have already mentioned. Each of those headlines, like (10), belongs to a generally reliable and pro-climate science news outlet. Some of these examples are "Brace Yourself: Top 10 Terrifying Impacts of Climate Change" (*abc News*), "7 Scary Facts About How Global Warming Is Scorching the United States" (*Mother Jones*), or "8 Charts That Show the Terrifying Reality of How Climate Change Is Affecting the U.S." (*salon*).

Obviously, in cases like (9), a clickbait headline alone is not enough to classify a news item as fake news. In fact, when defining institutional fake news, it is always necessary to assess the degree of the violation. While it is true that using clickbait headlines represents a deviation from standard journalistic writing practices and from the function a headline serves in this communicative context, it is also true that for the violation to qualify as fake news, it must be significant and must be evaluated on a case-by-case basis. For instance, if the NBC News article accompanying (9) is highly

informative and thus fully satisfies the cognitive need created by the use of a forward-referencing headline, the violation could be considered 'tolerable' and not significant.

On the other hand, when looking for news about climate change on blogs that typically spread misinformation on the topic – such as *Breitbart*, *Natural News*, or *InfoWars* – an opposite trend can be observed: the headlines used by these blogs are longer than average because they already include the entire argument they aim to promote. These headlines are, in fact, the opposite of forward-referencing titles, as they are redundant compared to the accompanying text. They condense all the informational content into the headline itself, allowing readers to avoid reading the text entirely and without making any effort. Here are some headlines taken from *Natural News*, an American blog that usually spreads disinformation on climate change:

- Climate change scientists LIE about everything by cherry picking different time windows for every study, graph or news announcement. (Huff 2019)
- Climate change data is based on FRAUD, and scientists around the world are pushing back against the narrative. (Johnson 2024)
- Corporate media relying on very recent history to claim global warming – but with a longer view of time, it's clear that human-caused climate change is a HOAX. (Huff 2024)

These headlines are not forward-referencing, for they do not invite the reader to read the accompanying article. They are autonomous texts.

One might argue that they violate norms of news publication, as the article text of each of them is redundant with respect to their headline or otherwise not sufficiently informative. Alternatively, they might be criticized for using capitalized words to grab attention, as in the case of (4) (p. 159). However, it would not be entirely accurate to describe these as cases of fake news resulting from clickbait headlines. It is, therefore, much more likely that these would be classified as instances of pure disinformation and/or as fake news but for other reasons (other than clickbait headlines).

It is interesting to note, however, that cases like these are difficult to address using a privative account of fake news. Consider, for example, the definition proposed by Lazer (2018), which is often adopted in the scientific literature on climate fake news. According to this definition, fake news is fabricated information that mimics news

media content in form but not in organizational process or intent. In the study by Lutzke et al. (2019), which we have already cited in §3.1, this very definition by Lazer and colleagues is applied to an article from *Natural News* featuring a headline stylistically similar to those mentioned above: "NASA confirms: sea levels have been FALLING across the planet for two years... media SILENT". This article is labelled as fake news by Lutzke and colleagues according to the privative definition of Lazer et al. However, it is far from clear that a headline like that shows similarity with the style and the appearance of real news headlines. On the contrary, the headline from *Natural News* is biased and unprofessional and it hardly complies with journalistic standards. This means that, although the authors of the study had a strong intuition that the *Natural News* article was a case of fake news, they relied on a definition of fake news – the one form Lazer and colleagues – that could not really classify the same article as fake news. This, in turn, indirectly suggests that a privative definition might fail to capture all those fake news instances that do not replicate the appearance of real news.

4.2 Deceptive headlines on climate change

In the previous chapter we said that deceptive headlines, as defined by Jodłowiec (2023), are manipulative headlines that create a mismatch between the meaning readers infer and the content of the associated article. While they do not necessarily contain outright falsehoods, they exploit semantic ambiguity to mislead audiences, encouraging incorrect interpretations. Unlike clickbait headlines, which tease readers by withholding information through forward-referencing (e.g., "You won't believe what happened"), deceptive headlines function as standalone, autonomous texts that appear definitive but mislead through suggestive phrasing. How do deceptive headlines apply to the context of climate change?

To identify deceptive headlines in the realm of climate change, one can look for patterns that create misleading or ambiguous impressions. These headlines often overstate findings, omit key context, or present speculative claims as definitive. Here is how climate deceptive headlines work. Consider the following headline from *Breitbart*, a far-right news website:

> 11) Climate Alarmists Finally Admit 'We Were Wrong About Global Warming'. (Delingpole 2017)

What does (11) communicate to its readers? According to the lexicon of climate sceptics climate alarmists are those people – be them scientists, politicians or common citizens – who highlight the relevance and the impact of climate change/global warming. However, this label often includes reputable climate scientists and/or epistemic authorities that embrace the scientific consensus on the anthropogenic climate change. So the explicature of (11) might be something like (11E):

> 11E) [People who highlight the relevance and the impact of global warming] admit [after some time that they] were wrong about global warming.

Notice that since "admit" is a factive verb, (11) carries the presupposition that at least some climate change advocates are wrong, and it could give rise to the unwarranted generalization that "Everything climate change advocates say is wrong"[85].

So (11) communicates a strong claim: people who highlight the relevance and impact of global warming – including reputable climate scientists and epistemic authorities – have admitted that what they said on global warming so far (or a part thereof) is wrong. The headline does not anticipate what or how much of what 'climate alarmists' say is wrong. But as we know by now, RT predicts that readers of (11) will process the headline on the assumption that it is newsworthy; that is, that it will satisfy their expectations of optimal relevance. Regarding the case of the Polish national team coach (p. 161), the audience of (11) will choose an interpretation of that "what" that leads to the greatest batch of cognitive effects for them. For instance, one plausible interpretation of this kind would be that so-called climate alarmists have admitted that no global warming is occurring.

But someone who reads the whole article finds out that this is not the right interpretation; indeed, the author of the article softens progressively the relevance of the interpretation at stake, to the point that the final claim is very weak when compared to the initial promise of relevance. Here is an extract of initial lines of the accompanying text of (11):

[85] Factive verbs – like "realize", "know", "confess" – are verbs that trigger the presupposition that the events or states conveyed by their sentential complements are true. For instance, the utterance "I know that the Italian gross domestic product has plummeted this year" triggers the presupposition that it is true that the Italian gross domestic product has plummeted this year.

> [...] a landmark paper, published in *Nature Geoscience*, which finally admits that the computer models have overstated the impact of carbon dioxide on climate and that the planet is warming more slowly than predicted.

Here we have the first weakening: what climate alarmists got wrong is not that a global warming is occurring, but that such warming, which is real and ongoing, is slower than predicted. This interpretation is still relevant, albeit not as relevant as the one envisaged before. But then we have another weakening:

> The paper – titled *Emission budgets and pathways consistent with limiting warming to 1.5 °C* – concedes that it is now almost impossible that the doomsday predictions made in the last IPCC Assessment Report of 1.5 degrees C warming above pre-industrial levels by 2022 will come true[86].

Now the interpretation of "what" in the headline refers to the fact that some specific predictions made in a single scientific document (the fifth IPCC assessment report of 2014) are probably wrong. This new interpretation is even weaker than the previous one in terms of potential cognitive effects[87].

Regarding cases of deceptive headlines analysed in the third chapter (§4.3), the article that accompanies (11) is still newsworthy to a certain degree, for the it may be relevant to know that a specific prevision of the sixth IPCC assessment report was wrong. The problem, again, is that this news content is not as newsworthy as its attached headline promised to be. Compare in fact (11) to its revised version, (11R):

> (11R) The concerned about climate finally admit 'We got the prediction of 1.5 degrees C warming above pre-industrial levels by 2022 wrong: it will take a few more years'.

The potential cognitive effects that readers may develop through processing (11R) are much less than those that may be developed through processing (11). So (11) is a case

[86] Pace Delingpole, today we know we have almost certainly passed the threshold of 1.5C in 2024 (Milman 2024), i.e. only two years later than predicted by the IPCC Assessment Report of 2014 (the one Delingpole is referring to).

[87] To be precise, it is true that the paper titled "Emission budgets and pathways consistent with limiting warming to 1.5 °C" (Millar et al. 2017) argues that models of the IPCC may have overestimated the immediate warming response to CO2, but it does not dismiss their utility. Instead, the paper calls for recalibrating carbon budgets based on updated observations.

of deceptive headline – and thus of climate fake news in my sense – although *it is not climate disinformation*. In fact, the claims in Delingpole (2017), if cleansed of exaggerations and subjective judgements, are fundamentally true: there exists a paper by Millar and colleagues (2017) that suggests that the climate models in the fifth IPCC assessment report might have overstated the warming rate relative to cumulative carbon dioxide emissions (although the implications that the *Breitbart*'s article derive from that single paper are unjustified, see footnote 86) .

4.3 Targeted climate fake news

In this section, I present two intriguing cases of targeted fake news on the topic of climate change. Targeted fake news, as explained in Chapter 3 (§4.4), are news items that violate the norms of production because, being constructed with a cover target, they exclude part of the news consumers from the informational exchange. Furthermore, the author of targeted fake news also breaches the norms of publication, as they do not take responsibility for the implications conveyed by the text in a context devoid of mutual adjustment (as defined in Chapter 3, p. 116).

In the case of fake news about climate change, we partially already know the potential target groups. These have been discussed both in the section on climate disinformation and, to some extent, in the review of existing literature on climate fake news. I am referring to those news consumers who share a clearly defined political orientation or ideology: libertarians or conservatives who perceive greenhouse gas emission reduction policies as a threat to the free market. The two examples I will illustrate below share this target group and are both taken from *Breitbart*. The first example has to do with the tremendous bushfires that hit Australia in 2019-2020.

Starting in June 2019 and peaking during the dry summer months, these fires ravaged vast areas across multiple states, particularly New South Wales and Victoria. The fires were fueled by prolonged drought, high temperatures, and strong winds, creating perfect conditions for uncontrollable blazes. Over 18 million hectares of land were burned, according to the UN, resulting in the destruction of thousands of homes and the tragic loss of 33 lives. The environmental impact was staggering, with numerous animals killed or displaced, pushing several species closer to extinction. Iconic wildlife, such as koalas, faced severe habitat loss. The fires also caused massive air quality issues, with

thick smoke blanketing cities like Sydney and Canberra, affecting millions of people and spreading as far as New Zealand.

In addition to the immediate destruction, the fires sparked a national and global conversation about climate change, as the crisis highlighted the increasing frequency and severity of extreme weather events. One of the clearest links to climate change was the record-breaking temperatures. Australia experienced its hottest year on record in 2019, with the average temperature 1.5 °C above the long-term average. December 2019 saw the highest national daily temperature ever recorded, reaching 41.9 °C. These extreme temperatures dried out vegetation and soil, creating a highly flammable environment. Prolonged drought, another factor linked to climate change, also played a critical role. Many parts of Australia experienced severe drought conditions leading up to the fires, which significantly reduced soil moisture and parched the vegetation. This made the landscape more susceptible to ignition and rapid fire spread.

In short, for many national and global commentators it was clear that climate change had definitely played a role in the devastation of Australian forests, even though it was not the primary cause of the bushfires. Nevertheless, amid the tragedy that hit Australian territories many right-wing blogs started pushing an antagonist narrative. These blogs, overtly sceptical of climate science, got across the message that Australian bushfires were due to the sole activity of arsonists. For instance, the conspiracy site *InfoWars*, that fuels climate change denialism, reported that nearly 200 people were arrested in Australia for deliberately starting bushfires. But as Palma (2020) explained at *Snopes*, a renowned fact-checking website, this report distorts the facts: most of those 200 people were just charged with deliberately setting fires and not arrested.

Furthermore, during the Black Summer – this is how media named the terrible season of bushfires in Australian in 2019-2020 – reports of arson were sometimes misrepresented or exaggerated. While deliberate ignitions did occur, the vast majority of fires were caused by natural factors, such as lightning strikes, and exacerbated by extreme weather conditions tied to climate change. But while *InfoWars*'s article was patently false, not so for one of *Breitbart* on the same topic, which came out with the following headline:

> 12) Police in Australia Begin Massive Criminal Investigation into Bushfire Arson. (Kent 2020)

(12) was paired with a text that does not even mention climate change as a contributing factor of the Australian bushfires. The text is just a collection of reports about arson accusations in Australia:

> On Monday, police in NSW reported they had already caught 24 alleged firebugs accused of deliberately lighting bushfires in a catastrophic season that has so far killed 18 people in the country's most populous state. [...] Two dozen people have allegedly risked prison sentences of up to 25 years by choosing to light bushfires. [...] A 16-year-old boy found to have started a fire that razed 14 homes in central Queensland and dealt with under the state's Youth Justice Act… (Kent 2020)

Notice that this article from Breitbart does not report false or misleading claims. It is true that individuals were charged of arson during the 2019-2020 Australian bushfires or that research has shown that many fires are caused by human activity, whether accidental or intentional. Overall, the statistics mentioned in the article are true too. For instance, the author observes that 85% of bushfires are caused by human activity. This is accurate, even though a bit out of context[88]. So we cannot straightforwardly say that (12) and its text are patently false.

However, if someone checks the comment section of the *Breitbart*'s article, they find out that the main message that the article gets across is one of climate denial. The vast majority of the over 4,800 comments come from users claiming that the article is a proof that climate change has nothing to do with the Black Summer. Let us emphasize this once again: the Breitbart article never mentions the term "climate change" (or "global warming"), not even to refute its role in the bushfires. This makes it particularly intriguing to understand how an article that never references a concept manages to communicate content about that very concept to thousands of users.

In my view, what is happening here is something very similar to the anti-vaccinationists' case (p. 167). The author of (12) wants to induce in the reader – and only in a specific category of readers, i.e. a target group – the recovery of a certain inference without taking the responsibility of that inference. The target group, here, is no more Italian anti-vaccinationists, but climate sceptics. Nevertheless, the strategy is

[88] This statistic includes both deliberate and accidental fires (e.g., campfires left unattended, sparks from machinery), many of which are unrelated to arson; however, the article focuses on arsons alone, giving the impression that this statistic refers only to deliberate bushfires.

the same: exploiting the difference of implicatures' accessibility of different readers of the same content. Here is how.

When we consider the cognitive processing of headlines like (12) (and its accompanying text), we need also to take into account the mutual cognitive environment between the author of the article and the readers. In the anti-vaccinationist case, my starting hypothesis is that the author of (12) and the news consumers, be them the target or not, share a set of assumptions about the events at stake. For instance, we may reasonably suppose that assumptions like "Bushfires are occurring in Australia", "There is an ongoing debate on the causes of Australian bushfires", and many others were all shared by the titlist and the news consumers. Therefore, the author of (12) could reasonably assume that an assumption like "Most experts agree that climate change facilitated the outbreak and the spreading of bushfires in Australia" is shared by people who read (12). And the same goes the other way: readers of (12) could reasonably assume that the author of (12) shared the same assumption (or the like). Therefore, we can say that the assumptions

(12A) There is an ongoing debate on the severity of Australian bushfires

(12B) Most experts agree that climate change facilitated the outbreak and the spreading of bushfires in Australia

are mutually manifest assumptions for both the author and all readers of (12).

Now considers libertarians and/or conservatives who are sceptical about the findings of climate science and oppose greenhouse gas emission reduction policies. Most of *Breitbart*'s readers and supporters share this ideology and worldview; that's why Breitbart acts as a representative media of far-right groups in US. Overall, this blog has shared over the years many articles which are overtly sceptical or even denialist concerning climate crisis. However, we may speculate that not all readers of Breitbart belong to the same ideological milieu: their audience is potentially wider. Libertarians and conservatives that do not trust climate science are just the target group of (12).

For them, (12A) and (12B) are manifest too. But as climate sceptical (or even denialist) they usually entertain other stronger assumptions about the topic of climate change that, as soon as they read (12), suddenly become *more manifest* to them. These

further and topic-specific assumptions may play a role in the recovery of specific weak implicatures from the reading of (12). Here is how.

Firstly, readers of the headline – be them the target group or not – reconstruct the explicature of (12). In this case the explicature of (12), call it (12E), might include contextual details that are inferentially added to the initial utterance. For instance, since the *Breitbart*'s article comments live on a specific news case, the outbreak of Australian bushfires, the following addition is part of the explicature:

(12E) [After the outbreak of major and unexpected bushfires in Australia,] Police in Australia Begin Massive Criminal Investigation into Bushfire Arson.

Then, readers may recover two implicit premises of (12):

($12IP_1$) Bushfires are either natural or human caused.

($12IP_2$) If the Australian police is investigating cases of arson, then the bushfires have probably been set out by arsonists alone.

From the combination of (12E), ($12IP_1$) and ($12IP_2$) readers might recover the implicated conclusion that

(12IC) Australian bushfires have probably been set out by arsonists alone[89].

This conclusion is strongly implicated because its recovery is needed in order to arrive at an interpretation that satisfies the readers expectations of optimal relevance.

Notice that the implicated conclusion (12IC) can be entertained by any news consumer, even non-target ones. The reason is that (12) strongly communicates it, and the author of the headline can hardly deny having intended this conclusion. However, people may not stop here in the recovery of implicatures.

Regarding the fictitious dialogue between Ken and Bev (p. 150), readers of (12) might go on to consider other possible implicatures for which the author offers gradually less evidence. For example, some of them might derive the implicature that Australian bushfires have been certainly set out by arsonists[90]. Other weak implicatures,

[89] I have chosen the headline of the *Breitbart*'s article for the sake of simplicity, but the analysis applies equally to the claims of the accompanying text.

[90] In fact, despite being more peremptory, this is a weaker implicature because the combination of (12E), (12IP1) and (12IP2) allows only for a conclusion that has a lighter epistemic modality. If the Australian police is investigating cases of arsons, it entails that police officials are not sure that bushfires

suggested by the word "massive", may be that the issue of bushfires is being treated as a high priority by the authorities, or that significant resources (manpower, funding, or technology) are being mobilized to carry on the investigation.

But climate sceptics – the target group – may entertain other weaker implicatures that are not equally accessible to non-target consumers. Unlike the case of Italian anti-vaccinationists, here the context expansion may be determined by the use of "criminal investigation" and "arson" and the overall emphasis of the article on deliberate human actions rather than climate factors. These linguistic markers, combined with the mutually manifest assumptions (12A) and (12B) and the ideologically-driven pre-existing assumptions on climate change, led climate sceptics to derive further implicatures concerning the role of climate factors.

For instance, a climate sceptic that reads (12) may derive a series of very specific weak implicatures, such as

($12IW_1$) Climate change has nothing to do with Australian bushfires.

($12IW_2$) Climate activists or media might be overlooking criminal activity in favour of a climate change narrative.

Unlike the case of Italian anti-vaccinationists, whose sneaky assumptions on vaccines became more manifest as soon as they encounter an expression with a specific encyclopaedic import, here the context expansion may be determined by the focus of (12) on arsons. This focus expands the context of interpretation of (12) so that it prioritizes human malice over natural or climate-related causes.

For a climate sceptic, assumptions like "The narrative on climate change is exaggerated" or "Bushfires are often misattributed to climate change when other factors, like arson, are to blame" are likely part of their cognitive environment. These assumptions are more easily accessible because, as we said in chapter 3 talking about accessibility in memory (p. 153), they may be frequently activated, well-practiced, or emotionally significant for a climate sceptic. Pre-existing worldview-related assumptions are often deeply ingrained, having been strengthened by repeated

---

were caused by arsonists; they are assessing this hypothesis, maybe they have found it convincing, but they are not certain about it (otherwise the investigation would have been already concluded).

activation through personal experiences and social reinforcement (especially through online filter bubbles and echo chambers)[91].

Therefore, headlines – and news articles in general – like (12) manage to induce the recovery of a range of a climate-related weaker implicatures in a certain portion of the audience, but without explicitly doing so, because (12) neither mention the expression "climate change" and related ones nor pronounce on the issue of the contribution of climate factors. Therefore, (12) is targeted climate fake news: it does not present overtly false claims, as sheer climate misinformation, but its focus on arsons may strengthen the narrative of the irrelevance of climate factors in some portion of the audience.

A second example of targeted climate fake news comes again from *Breitbart*. It is very similar to the previous one in its functioning, but it is even more striking for its apparent neutrality. On 22 February 2023, the blog posts a news item with the following headlines:

> (13) San Francisco Braces for First Snow Since 1976 as More Winter Storms Arrive. (Pollak 2023)

As for (12), both the title and the accompanying text of (13) made no apparent reference to climate change and related discussions. A unique refence could be found in a separate section of suggested articles on the bottom of the page, where readers can choose a selection of articles divided according to the topic. Among them, there is one denoted through the topic category of "climate change". Apart from that, the main article just highlights the fact that San Francisco prepared for a rare snowfall in February 2023, marking the first potential snow in the city since 1976.

However, a cursory glance at the comment section reveals a flood of remarks mocking climate science. Examples include statements like this:

> The tendency of so-called scientists to measure world events on an insignificant timeline like a human being's lifespan is the reason so much of "settled" climate science is just garbage…. We live on a planetary mass that's billions of years old and I can guarantee that there have been no two 24 hour days in it's history with "identical" climate on the planet in the entire timeline… yet the elites want to instill panic with something no one has any control over!!

[91] For more information on climate echo cambers and how they play a role in strengthening ideological beliefs and assumptions on climate change, see Walter et al. (2017); Jasny and Fisher (2019); van Eck et al. (2021).

Current human progress is the result of planetary conditions not the cause of them! Period! (Anonymous user)

(13) has prompted many users to comment with sceptical narratives on climate change, despite the content of (13) has no meaningful relation with that topic (apart from being a weather forecast). Besides, unlike (12), (13) has not been published in the middle of a vivid global debate about bushfires or other extreme weather events: it just came out of thin air. Once again, we may use the framework of Relevance Theory to dissect the whole inferential process.

Assuming that the explicature of (13) is very similar to (13) itself, the first implicated premise to occur to the readers of the headline will likely be:

(13IP) If San Francisco does not experience snowfall since 1976, then the incoming snowfall is a rare event.

Thus, from the combination of the explicature of (13) and (13IP), readers will interpret (13) as conveying the following implicated conclusion:

(13IC) The incoming snowfall in San Francisco is a rare event.

However, *Breitbart*'s target group – predisposed to climate scepticism – draws on their pre-existing cognitive environment to derive further implicit meanings. So for *Breitbart*'s audience, especially its core libertarian and conservative readers, additional assumptions like

(13C) Unusual cold weather contradicts narratives of global warming

may be more manifest[92]. These deeply ingrained beliefs are ideologically reinforced through *Breitbart*'s content ecosystem, where scepticism of climate science is common. Assumptions like (13C), in conjunction with (13IP), may lead the target group to infer further implicated conclusion such as:

(13IW1): Rare cold events, like San Francisco's snowfall, disprove the theory of climate change.

[92] Notice that this specific assumption is patently false: just as a cold day during winter does not mean winter is gone, in the same way an unusually cold event does not negate the broader trend of a warming planet.

(13IW2): Climate narratives overlook natural variability to promote human-caused warming.

But from the point of view of RT, (13IW1) and (13IW2) are very weak implicatures, because nothing in (13) or in the accompanying article makes any reference, even indirectly, to the content of these additional implicated conclusions.

These inferences are accessible to the target groups due to ideological pre-existing beliefs and the echo chambers built around blogs like *Breitbart*, where information supporting scepticism is amplified. Such factors explain why weak implicatures like (13IW1) and (13IW2) feel convincing to this audience. Interpreting the article through their worldview requires minimal cognitive effort while reinforcing their existing beliefs, offering what Relevance Theory calls optimal relevance. For non-sceptical readers, the same headline may prompt entirely different interpretations. They will probably focus on the rarity of snowfall without connecting it to climate change debates. This highlights how ideological predispositions mediate access to implicatures, and why headlines such as (13) are targeted fake news.

These cases of targeted climate fake news we have examined are very interesting because they entirely or partially escape traditional fake news accounts. For instance, a hybrid account of fake news can hardly claim that (13) is fake news from its own perspective. In fact, even assuming there is a bad epistemic attitude behind the publication of (13) (which would not be too surprising), it remains true that neither the headline nor the accompanying article contains false or misleading information. Yet, we discover from the comments section that, for a significant number of people, (13) and the related article conveyed false content about climate change, even though the discussion on the climate crisis was not explicitly mentioned.

It is clear, therefore, that if there is something wrong with (13), it cannot be attributed to its propositional or informational content. It is not, in other words, about communicating falsehoods but rather about an improper way of presenting information which, when threw to a specific audience in a particular context (that of a blog openly sceptical about climate change), ends up inducing readers to infer something that cannot easily be attributed to the article's author. In my view, this represents a strategic use of the violation of norms of news publication, which require, among other things, avoiding leaving room for the calculation of implicit meanings that the article does not intend to

convey. It is a particular type of violation because the context of journalistic communication is itself equally peculiar, as we explained in the previous chapter when discussing the absence of a context of mutual adjustment (p. 116).

Conclusion

This chapter has delved into the issue of climate fake news, situating it within the broader context of disinformation and exploring its unique characteristics. By adopting an institutional framework, it has argued that fake news differs from other information pathologies due to its violation of journalistic norms. While climate disinformation has received considerable academic attention, climate fake news as a distinct phenomenon remains underexplored, prompting the need for a more nuanced understanding.

The discussion began by defining climate change and contextualizing it within its scientific and societal implications. Climate change represents a global challenge, deeply tied to political, economic, and ideological divisions, which makes it a fertile ground for the propagation of misleading contents. Central to the chapter's analysis was the distinction between climate disinformation and climate fake news. Unlike disinformation campaigns orchestrated by specific interest groups and aimed at casting doubt on climate science through carefully crafted narratives, the existing scientific literature, albeit scant, tells us that climate fake news frequently operates at the intersection of media sensationalism and ideological partisanship, often leveraging emotional appeals and social media dynamics to amplify its reach.

In the literature, climate fake news has often been categorized according to a privative criterion: it is news that mimics the appearance of real news but fails to meet the journalistic standards, and whose content deals with the topics of climate change and/or sustainability. Sometimes, albeit more rarely, the category of climate fake news resonates more with a conservative criterion, which requires news to be false or misleading and intentionally so. In this chapter, I suggested adopting an institutional framework for climate fake news. According to that, climate fake news is news that deals with the topic of sustainability and/or climate change and which significantly violates the norms of journalism (or, equivalently, does not live up to a sufficient degree to realistic expectations towards the products of a journalistic source).

Through the analysis of real cases that mimics the scheme of the third chapter, here I illustrated how climate fake news functions in practice. Firstly, I claimed that, paradoxically, the risk of finding climate fake news in the form of clickbait articles is higher for renowned newspaper than some blogs which actively spread climate disinformation[93]. In fact, many sources of disinformation on climate change openly endorse an antagonist attitude towards climate science and frequently use headlines that are autonomous texts and presents an explicit sceptical thesis.

Afterwards, I presented a case of deceptive headline on climate change. Unlike the examples in the previous chapter, for obvious reasons climate fake news does not adopt death hoaxes to be deceptive (in the sense envisaged by Jodłowiec). Instead, I brought up the case of a headline – "Climate Alarmists Finally Admit 'We Were Wrong About Global Warming'" – that creates a semantic mismatch between the expectations it sets and the actual content of its accompanying article. The use of loaded terms such as "alarmists" immediately primes the audience for a narrative of scientific fallibility, inviting interpretations that maximize cognitive effects. The misleading emphasis on "finally admit" suggests a sweeping repudiation of climate science, yet the article ultimately focuses on a narrow issue of modelling inaccuracies related to the rate of warming.

Finally, I have brought up two cases of targeted climate fake news. The analysis of *Breitbart* headlines on Australian bushfires and rare snowfall in San Francisco highlighted how these articles strategically omit explicit claims about climate change while implicitly inviting readers to draw conclusions aligned with climate scepticism. Using Relevance Theory, the chapter demonstrated how weak implicatures allow such articles to subtly reinforce ideological divides without overtly engaging with the scientific discourse on climate change. These findings suggest that climate fake news sometimes operates through a deliberate selection of target groups, tailoring its content to resonate with ideologically predisposed news consumers. For sceptical readers, ideological biases and echo chambers increase the accessibility of weak implicatures, enabling the recovery of climate-denialist narratives even in articles that do not explicitly mention climate change. This mechanism, facilitated by targeted headlines,

[93] It always depends on whether the accompanying article is sufficiently informative – that is, if the cognitive reward of its reading lives up to the expectations.

underscores how fake news perpetuates polarization, reinforcing scepticism among its intended audience while avoiding direct accountability for its implications.

In conclusion, this chapter has argued that climate fake news, while less researched than disinformation, represents a distinct and insidious challenge to public understanding of climate change. Its reliance on weak implicatures and targeted framing allows it to exploit ideological divisions, undermining trust in science and delaying collective action. The adoption of an institutional account may enable us to diagnose new and interesting cases of climate fake news that are not equally captured by conservative and privative perspectives.

Concluding remarks

In this dissertation, I sought to redefine the philosophical notion of fake news to address the challenge of conceptual necessity raised by Habgood-Coote in 2019 and to more effectively analyze cases of fake news about climate change. The first objective was pursued across the first three chapters, while the second was reserved for the final chapter.

The first chapter laid the historical and definitional groundwork, tracing the evolution of the term “fake news” from its early conceptual usages to its contemporary manifestations as both a journalistic and philosophical concern. By reviewing its historical roots and epistemological debates, the thesis identified two prevailing models: the conservative view, which treats fake news as a subset of news, and the privative view, which denies its inclusion within the domain of news. The chapter highlighted the instability of the term, marked by shifting meanings, its weaponization in political discourse, and its redundancy due to overlaps with existing concepts. These issues underscored the challenge of conceptual necessity raised by Habgood-Coote.

In short, most definitions of fake news in the literature – primarily those aligned with the conservative view – risk being unnecessary and redundant because other, more precise notions (e.g., misinformation, disinformation, propaganda, or lying) already perform the same explanatory work. Privative definitions, while not directly subject to Habgood-Coote’s critique, either suffer from similar issues on a different level or introduce new limitations.

To propose a novel approach, the second chapter observed that different definitions of fake news are grounded in different conceptions of what constitutes news. Conservative accounts, which dominate the literature, typically treat news as a form of information, such as information reported through a news format. In contrast, privative accounts view news as a status that information may or may not attain, depending on its history of production. According to this perspective, the correct history of production endows information with specific value derived from its enhanced epistemic quality.

Fake news, then, is a story that merely appears to be news but lacks the proper history of production and, consequently, the epistemic value associated with genuine news.

While I regard the conservative approach as patently naïve, I argued that the privative conception, though incomplete, is closer to the mark. News is more of a status than an object, and we should avoid conflating news with information, as is common in everyday conversation. Expanding on the intuition of the privative approach, I proposed a more fine-grained view that conceptualizes news as an institutional entity – a perspective inspired by John Searle's theory of institutional facts. News, in this account, is a status function bestowed on information by journalistic sources, conferring a set of rights and obligations on both journalists and readers. This institutional perspective resolves some of the conceptual ambiguities of the privative approach while emphasizing the normative expectations that distinguish news from other forms of communication.

Building on this institutional account, the third chapter proposed a novel definition of fake news: "news that does not deserve its status." This definition bifurcates into two complementary interpretations. The first (FN_A) focuses on the violation of journalistic norms, encompassing breaches such as failures in verification, transparency, and many others. The second (FN_B) addresses the failure of news to meet the realistic expectations of its consumers, which are themselves informed by the professional standards of journalism. These perspectives are interrelated: violations of journalistic norms (FN_A) lead to the disappointment of consumer expectations (FN_B). Together, they illuminate the dual dimensions of fake news: as a normative breach from the perspective of producers and as a failure to meet expectations from the perspective of consumers. This dual approach reconciles the views of journalists and audiences, integrating them into a cohesive framework that underscores their interdependence.

The institutional account addresses the challenge of conceptual necessity by disentangling fake news from notions of truth and intentionality. As argued throughout the dissertation, fake news can be true (e.g., accidentally) and involuntary (e.g., due to severe malpractice). By anchoring the concept in journalistic norms and expectations, the thesis establishes fake news as a distinct and meaningful term within the discourse on information pathologies. Terms like "disinformation" address the truthfulness of a

news item but fail to capture the normative breach associated with violations of journalistic rules.

With this framework, the third chapter also explored additional cases of fake news. Norms of journalism, in fact, regulate the phase of news production as well as the one of news publication. This diversification opens the gates to many kinds of violations, which in turn entails more and new cases of fake news that may have not been captured by already existing accounts on the market. For instance, we have seen that the violation of norms of news publication often involves the use of pragmatic strategies that, starting from the headline of a news item, may fool the reader into thinking that a story is very newsworthy when actually is not or exploit audience stratification to get across undercover very specific implicatures. This discussion segued into the dissertation's final contribution: the investigation of fake news in the context of climate change.

The scientific literature on climate disinformation has grown rapidly over the past decade. Given the polarizing nature of the climate crisis, many scholars and commentators have speculated about the role of fake news in fostering confusion and skepticism around climate science. However, studies claiming to address climate fake news often focus instead on climate disinformation. This conflation occurs for two main reasons:

1. The term fake news is frequently used interchangeably with mis/disinformation in journalistic contexts, reinforcing the need for conceptual clarity;
2. The term fake news has been politicized, rendering it less serviceable in scientific discourse, which seeks to avoid politically charged terminology.

As a result, research specifically focused on climate fake news is limited compared to studies on climate disinformation, despite overlaps in their mechanisms of propagation and target audiences.

Existing studies on climate fake news predominantly adopt a privative framework. However, if we define climate fake news as "any news item about climate change that significantly violates journalistic norms," we can uncover cases that would otherwise go undetected. This is the kind of operation that I pursued in the fourth chapter. Many such cases involve deceptive headlines. For instance, an article might promise a significant debunking of climate science findings but, upon closer reading, offers a much weaker claim. These cases qualify as climate fake news not because their content is false or

misleading but because they fall short of journalistic expectations of newsworthiness, creating inflated interpretations in readers.

More sophisticated cases include targeted fake news. For example, articles about snow in San Francisco may appear to focus on benign topics like weather forecasts but are crafted for ideologically predisposed audiences, such as climate skeptics. Such readers interpret these articles through erroneous assumptions, such as believing unusual cold spells disprove global warming. These bait articles exploit audience predispositions, leaving readers to infer incorrect conclusions without the author's explicit endorsement. These cases cannot be classified as mis/disinformation in the classical sense (e.g., false or misleading information spread with the intent to deceive). Instead, they represent a strategy to reinforce pre-existing beliefs or make them more salient for specific goals. This strategy involves a different kind of intentionality that warrants further investigation.

A promising avenue for future research would be to explore the diverse notions of intentionality involved in disinformation phenomena. Starting with the targeted fake news cases analyzed in this dissertation, one could investigate how existing frameworks of miscommunication fail to account for audience predispositions that facilitate the spread of reinforcing messages. While Relevance Theory offers partial insights through concepts like mutual cognitive environment and encyclopedic knowledge, it falls short of fully explaining these dynamics. The next challenge, then, is to extend and refine analytical tools for studying communication to better address these emerging phenomena.

References

«Climate Change Catastrophes Ahead? Here's How We Can Prepare», *NBC News*, https://www.nbcnews.com/science/environment/climate-change-catastrophes-ahead-heres-how-we-can-prepare-n98066 (last access: 12/12/2024).

Adler J. E., 2002, *Belief 's Own Ethics*, Cambridge, MIT Press.

Al-Rawi A., O'Keefe D., Kane O., Bizimana A.-J., 2021, «Twitter's Fake News Discourses Around Climate Change and Global Warming», Frontiers in Communication, 6, 729818, pp. 1-9.

Allen M., Gillett N., Kettleborough J., et al., 2006, «Quantifying anthropogenic influence on recent near-surface temperature change», Surveys in Geophysics, 27, pp. 491–544.

Althaus S. L., Edy J. A., Phalen P. F., 2001, «Using substitutes for full-text news stories in content analysis: which text is best?», *American Journal of Political Science*, 45, 3, pp. 707–723.

Anderau G., 2021, «Defining Fake News», *Kriterion – J. Philos.*, 35, 3, pp. 197-215.

Andrew B. C., 2007, «Media-generated Shortcuts: Do Newspaper Headlines Present Another Roadblock for Low-information Rationality?», *The International Journal of Press/Politics*, 12, 2, pp. 24-43.

Asr F. T., Taboada M., 2019, *Big data and quality data for fake news and misinformation detection*, *Big Data & Society*, 6, 1, pp. 1-14.

Assimakopoulos S., 2017, «Context in Relevance Theory», in Blochowiak J., Grisot C., Durrleman S., Laenzlinger C. (eds.), *Formal Models in the Study of Language*, New York, Springer, pp. 221-242.

Axt J. R., Landau M. J., Kay A. C., 2020, «The Psychological Appeal of Fake-News Attributions», *Psychological Science*, 31, 7, pp. 848-857.

Barsalou L. W., 1987, «The instability of graded structure: Implications for the nature of concepts», in Neisser U. (ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization*, Cambridge, Cambridge University Press, pp. 101-140.

Barsalou L. W., 1989, «Intraconcept similarity and its implications for interconcept similarity», in Vosniadou S., Ortony A. (eds.), *Similarity and analogical reasoning*, Cambridge, Cambridge University Press, pp. 76-121.

Beaujon, 2019, «Trump Claims He Invented the Term “Fake News”—Here’ s an Interview With the Guy Who Actually Helped Popularize It», *Washingtonian*, https://www.washingtonian.com/2019/10/02/trump-claims-he-invented-the-term-fake-news-an-interview-with-the-guy-who-actually-helped-popularize-it/ (last access: 31/07/2024)

Bergamaschi Ganapini M., 2021, «The signaling function of sharing fake stories», *Mind & Language*, pp. 1-17.

Bernecker S., Flowerree A., Grundmann T., 2021, *The epistemology of fake news*, Oxford, Oxford University Press, pp. 310-333.

Björnberg K. E., Karlsson M., Gilek M., Hansson S. O., 2017, «Climate and environmental science denial: A review of the scientific literature published in 1990–2015», Journal of Cleaner Production, 167, pp. 229-241.

Blom N. J., Hansen R. K., 2015, «Click bait: Forward-reference as lure in online news headlines», *Journal of Pragmatics*, 76, pp. 87-100.

Bolin J. L., Hamilton L. C., 2018, «The News You Choose: News media preferences amplify views on climate change», Environmental Politics, 27, 3, pp. 455–476.

Bondielli A., Marcelloni F., 2019, «A survey on fake news and rumour detection techniques», *Information Sciences*, 497, pp. 38-55.

Boykoff M. T., 2008, «Lost in translation? United States television news coverage of anthropogenic climate change, 1995–2004», Climatic Change, 86, pp. 1–11.

Brontsema R., 2004, «A Queer Revolution: Reconceptualizing the Debate Over Linguistic Reclamation», *Colorado Research in Linguistics*, 17, pp. 1-17.

Brown E., 2019, «Fake News and Conceptual Ethics», *Journal of Ethics and Social Philosophy*, 16, 2, pp. 144-154.

Brüggemann M., Engesser S., 2017, «Beyond false balance: How interpretive journalism shapes media coverage of climate change», Global Environmental Change, 42, pp. 58–67.

Burge T., 1979, «Individualism and the mental», *Midwest Studies in Philosophy*, 4, 1, pp. 73-122

Cacioppo J. T., Gardner W. L., Berntson G. G., 1999, «The affect system has parallel and integrative processing components: Form follows function», Journal of Personality and Social Psychology, 76, 5, pp. 839–855.

Camp E., 2018, «Insinuation, Common Ground, and the Conversational Record», in Fogal D., Harris W. D., Moss M. (eds.), *New Work on Speech Acts*, Oxford University Press, Oxford, pp. 40-66.

Cappelle B., Denis P., Keller M., 2018, «Facing the facts of fake: A distributional semantics and corpus annotation approach», *Yearbook of the German Cognitive Linguistics Association*, 6, 1, pp. 9-42.

Carrasco-Farré C., 2022, «The fingerprints of misinformation: how deceptive content differs from reliable sources in terms of cognitive effort and appeal to emotions», Humanities and Social Sciencies Communications, 9, 162, pp. 1-18.

Carson T., 2006, «The Definition of Lying», *Noûs*, 40, 2, pp. 284-306.

Carson T., 2010, *Lying and Deception*, Oxford, Oxford University Press.

Carston R., 2002, *Thoughts and Utterances: The Pragmatics of Explicit Communication*, Oxford, Blackwell.

Cevolani G., 2011, «Strongly Semantic Information and Verisimilitude», *Etica and Politica / Ethics and Politics*, 13, 2, pp. 159-179.

Chakraborty A., Paranjape B., Kakarla S. and Ganguly N., 2016, «Stop Clickbait: Detecting and preventing clickbaits in online news media», *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pp. 9-16.

Chakraborty A., Paranjape B., Kakarla S., Ganguly N., 2016, «Stop clickbait: Detecting and preventing clickbaits in online news media», *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pp. 9 – 16.

Chalmers D. J., 2020, «What is conceptual engineering and what should it be?», *Inquiry*, pp. 1-18.

Charnley V. M., 1966, *Reporting*, New York, Holt, Rinehart and Winston.

Clark B., 2013, *Relevance Theory*, Cambridge, Cambridge University Press.

Clarke C. E., Hart P.S., Schuldt J.P, Evensen D., Boudet H., Jacquet J., Stedman R., 2015, «Public opinion on energy development: the interplay of issue framing, top-of-mind associations, and political ideology», *Energy Policy*, 81, pp. 131-140.

Cook J., Lewandowsky S., Ecker U. K. H., 2017, «Neutralizing misinformation through inoculation: Exposing misleading argumentation techniques reduces their influence», PLOS ONE, 12, 5, e0175799, pp. 1-21.

Cook J., Nuccitelli D., Green S. A., Richardson M., Winkler B., et al., 2013, «Quantifying the consensus on anthropogenic global warming in the scientific literature», Environmental Research Letters, 8, 024024, pp. 1-7.

Cook J., Oreskes N., Doran P. T., Anderegg W. R. L., Verheggen B., et al., 2016, «Consensus on consensus: a synthesis of consensus estimates on human-caused global warming», Environmental Research Letters, 11, 048002, pp. 1-7.

Croce M., Piazza T., (2022), *Che cosa sono le fake news*, Roma, Carocci.

Croce M., Piazza T., 2021, «Misinformation and Intentional Deception. A Novel Account of Fake News», in Snow N., Vaccarezza S. M. (eds.), *Virtues, Democracy, and Online Media: Ethical and Epistemic Issues*, New York, Routledge, pp. 49-63.

Croce M., Piazza T., 2021,«Consuming Fake News: Can We Do Any Better?», *Social Epistemology*, pp. 1-10.

D'Alfonso S, 2011, «On Quantifying Semantic Information», *Information*, 2, 1, pp. 61-101.

Daron J. D., Stainforth D., 2013, «On predicting climate under climate change», Environmental Research Letters, 8, pp. 1–8.

Del Pinal G., 2015, «Dual Content Semantics, privative adjectives, and dynamic compositionality», *Semantics & Pragmatics*, 8, 7, pp. 1-53.

Delingpole J., 2017, «Delingpole: Climate Alarmists Finally Admit 'We Were Wrong About Global Warming'», *Breitbart*, https://www.breitbart.com/politics/2017/09/19/delingpole-climate-alarmists-finally-admit-we-were-wrong-about-global-warming/ (last access: 12/12/2024).

Dinneen J.D., Brauner C., 2015, «Practical and Philosophical Considerations for Defining Information as Well-Formed, Meaningful Data in the Information Sciences», *Library Trends*, 63, 3, pp. 378-400.

Dominguez Armas A., Soria Ruiz A., 2021, «Provocative insinuations», *Daimon. Revista Internacional de Filosofía*, 84, pp. 63-80.

Doney S. C., et al., 2012, «Climate change impacts on marine ecosystems», *Annual Review of Marine Science*, 4, 1, pp. 11–37.

Dowd K., 2001, «The Emergence of Fiat Money: A Reconsideration», *Cato Journal*, 20, 3, pp. 467-476.

Dretske F., 1981, *Knowledge and the Flow of Information*, Cambridge, MIT Press.

Dretske F., 1983, «Précis of Knowledge and the Flow of Information», *The Behavioral and Brain Sciences*, 6, pp. 55-90.

Dretske I. F., 1981, *Knowledge and the Flow of Information*, Stanford, MIT Press.

Dretske I. F., 1983, «Précis of Knowledge and the Flow of Information», *Behavioral and Brain Sciences*, 6, 1, pp. 55-63.

Drummond C., Siegrist M., Arvai J., 2020, «Limited effects of exposure to fake news about climate change», Environmental Research Communications, 2, 081003, pp. 1-11.

Dymnikov V., Gritsoun A., 2001, «Atmospheric model attractors: Chaoticity, quasi-regularity and sensitivity», Nonlinear Processes in Geophysics, 8, pp. 201–208.

Egelhofer L. J., Lecheler S., 2019, «Fake news as a two-dimensional phenomenon: a framework and research agenda», Vol. 43: *Annals of the International Communication Association*, 2, pp. 97-116.

Elsasser S. W., Dunlap R. E., 2013, «Leading voices in the denier choir: Conservative columnists' dismissal of global warming and denigration of climate science», American Behavioral Scientist, 57, 6, pp. 754–776.

Fallis D., 2015, «What is Disinformation?», *Library Trends*, 63, 3, pp. 401-426.

Fallis D., 2016, «Mis- and dis-information», in Floridi L. (ed.), *The Routledge Handbook of Philosophy of Information*, New York, Routledge, pp. 332-346.

Fallis D., Mathiesen K., 2019, «Fake news is counterfeit news», *Inquiry: An Interdisciplinary Journal of Philosophy*, pp. 1-20.

Fallis D., Stokke A., 2017, «Bullshitting, Lying, and Indifference toward Truth», *Ergo*, 4, pp. 277-309.

Farrell J., 2016, «Corporate funding and ideological polarization about climate change», Proceedings of the National Academy of Sciences of the United States of America, 113, 1, pp. 92–97.

Farrell J., 2016, «Network structure and influence of the climate change counter-movement», Nature Climate Change, 6, pp. 370–374.

Ferragina E., Quagliarotti D. A., 2019, «Climate change and environmentally induced migration», in Ferragina E. (ed.), *Mediterranean Migration and the Labour Markets*, New York, Routledge, pp. 171-197.

Floridi L., 2005, «Is Semantic Information Meaningful Data?», *Philosophy and Phenomological Research*, 70, 2, pp. 351–370.

Floridi L., 2005a, «Is Semantic Information Meaningful Data?», *Philosophy and Phenomological Research*, LXX, 2, pp. 351-370.

Floridi L., 2010, *Information: A Very Short Introduction*, Oxford, Oxford University Press.

Floridi L., 2011, *The Philosophy of Information*, Oxford, Oxford University Press.

Fortunato S., 2010, «Community detection in graphs», Physics Report, 486, 3-5, pp. 75-174.

Fox J. C., 1983, Information and Misinformation: An Investigation of the Notions of Information, Misinformation, Informing, and Misinforming, London, Bloomsbury Academic.

Frankfurt H. G., 2005, *On Bullshit*, Princeton, Princeton University Press.

Frege G., 1892/1948, «Sense and Reference», *The Philosophical Review*, 57, 3, pp. 209-230, https://doi.org/10.2307/2181485.

Fricker E., 2012, «Stating and Insinuating», *Proceedings of the Aristotelian Society*: Supplementary Volumes, 86, pp. 61-94.

Frigg R., Thompson E., Werndl C., 2015, «Philosophy of Climate Science Part I: Observing Climate Change», Philosophy Compass, 10, pp. 953–964.

Gabielkov M., Ramachandram A., Chaintreau A., Legout A., 2016, «Social Clicks: What and Who Gets Read on Twitter?», *Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science*.

Galeotti E. A., Meini C., 2022, «Scientific Misinformation and Fake News: A Blurred Boundary», *Social Epistemology*, pp. 1-16.

Gelfert A., 2018, «Fake news: a definition», *Informal Logic*, 38, 1, pp. 84–117.

Gelfert A., 2021, «Fake News, False Beliefs, and the Fallible Art of Knowledge Maintenance», in Bernecker S., Flowerree A., Grundmann T. (eds.), *The epistemology of fake news*, Oxford, Oxford University Press, pp. 310-333.

Gentzkow M., Shapiro M. J., Stone F. D., 2015, «Media Bias in the Marketplace: Theory», in Anderson P. S., Waldfogel J., Strömberg D. (eds.), *Handbook of Media Economics*, North-Holland, Elsevier, pp. 623-645.

Gettier E., 1963, «Is Justified True Belief Knowledge?», Analysis, 23, 6, pp.121-123.

Gettings M., 2007, «The Fake, the False, and the Fictional: The Daily Show as News Source», in Holt j. (ed.), *The Daily Show and Philosophy*, Oxford, Blackwell, pp. 16 – 27.

Goldhamer M., 2022, «'Died Suddenly' film amplifies false Covid-19 vaccine claims», *AFP Fact Check*, Online: https://factcheck.afp.com/about-afp (last access: 10/10/2024).

Grice H. P., 1975/1989, «Logic and Conversation», in Grice H. P., *Studies in the way of words*, Cambridge, Harvard University Press, pp. 22-40.

Grieve J., Woodfield H., 2023, *The Language of Fake News*, Cambridge, Cambridge University Press.

Grundmann T., 2020, «Fake news: the case for a purely consumer-oriented explication», *Inquiry*, 66, 10, pp. 1758-1772.

Gustafson A., Rosenthal S. A., Ballew M. T., et al., 2019, «The development of partisan polarization over the Green New Deal», Nature Climate Change, 9, pp. 940–944.

Habgood-Coote J., 2019, «Stop Talking About Fake News!», *Inquiry*, 62, pp. 1033–1065.

Habgood-Coote J., 2022, «Fake news, conceptual engineering, and linguistic resistance: reply to Pepp, Michaelson and Sterken, and Brown», Inquiry, 65, 4, pp. 1–29.

Hallegatte S., et al., 2016, *Shock waves: Managing the impacts of climate change on poverty*, Washington D.C., World Bank Group.

Hallin C. D., Mancini P., 2004, *Comparing Media Systems: Three Models of Media and Politics*, Cambridge, Cambridge University Press.

Hampton M., 2004, *Visions of the Press in Britain, 1850-1950*, Champaign, University of Illinois Press

Hannon M., 2021, «Disagreement or Badmouthing? The Role of Expressive Discourse in Politics», in Edenberg E., and Hannon M. (eds), *Political Epistemology*, New York, Oxford University Press, pp. 297-317.

Haslanger S., 2000, «Gender and race: (What) are they? (What) do we want them to be?», *Noûs* 34, 1, pp. 31-55.

Haslanger S., 2020, «Going On, Not in the Same Way», in Burgess Al, Cappelen H., Plunkett D. (eds), *Conceptual Engineering and Conceptual Ethics*, Oxford, Oxford University Press, pp. 230-260.

Henderson R., Mccready E., 2019, «Dogwhistles and the At-Issue/Non-At-Issue Distinction», in Gutzmann D. and Turgay K. (eds.), *Secondary Content: The Semantics and Pragmatics of Side Issues*, Leida, Brill, pp. 222-45.

Hoegh-Guldberg O., et al., 2018, «Impacts of 1.5°C of global warming on natural and human systems», in Global Warming of 1.5°C, Intergovernmental Panel on Climate Change (IPCC).

Horne D. B., Adali S., 2017, «This Just In: Fake News Packs a Lot in Title, Uses Simpler, Repetitive Content in Text Body, More Similar to Satire than Real News», *ArXiv*.

Hornsey M., Harris E., Bain P., Fielding K. S., 2016, «Meta-analyses of the determinants and outcomes of belief in climate change», Nature Climate Change, 6, pp. 622–626.

Huff E., 2019, «Climate change scientists LIE about everything by cherry picking different time windows for every study, graph or news announcement», *Natural News*, https://www.naturalnews.com/2019-10-14-climate-change-scientists-lie-about-everything-cherry-picking-data.html (last access: 12/12/2024).

Huff E., 2024, «Corporate media relying on very recent history to claim global warming – but with a longer view of time, it's clear that human-caused climate change is a HOAX», *Natural News*, https://www.naturalnews.com/2024-09-26-media-history-global-warming-deception-data-study.html (last acess: 12/12/2024).

Ifantidou E., 2009, «Newspaper headlines and relevance: Ad hoc concepts in ad hoc contexts», *Journal of Pragmatics*, 41, pp. 699-720.

IPCC, 2007, Climate Change 2007: Synthesis Report. Contribution of Working Groups I, II and III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, Pachauri R. K., Reisinger A. (eds.), Geneva, IPCC, pp. 104.

IPCC, 2014, Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, Pachauri R. K., Meyer L. A. (eds.), Geneva, IPCC, pp. 151.

IPCC, 2019, «Annex I: Glossary», in Shukla P. R., Skea J., Calvo Buendia E., Masson-Delmotte V., Pörtner H.-O., Roberts D. C., Zhai P., Slade R., Connors S., van Diemen R., Ferrat M., Haughey E., Luz S., Neogi S., Pathak M., Petzold J., Portugal Pereira J., Vyas P., Huntley E., Kissick K., Belkacemi M., Malley J. (eds.), Climate Change and Land: an IPCC special report on climate change, desertification, land degradation, sustainable land management, food security, and greenhouse gas fluxes in terrestrial ecosystems, Cambridge, Cambridge University Press.

IPCC, 2023, «Sections», in Lee H., Romero J. (eds.), Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change, Geneva, IPCC, pp. 35–115.

Jacomy M., Venturini T., Heymann S., Bastian M., 2014, «ForceAtlas2, a Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software», *PLoS ONE*, 9, 6: e98679.

Jasny L., Fisher D., 2019, «Echo chambers in climate science», Environmental Research Communications, 1, 10, pp. 1-8.

Jaster R., Lanius D., 2018, «What is Fake News?», *Versus*, 2, 127, pp. 207–227.

Jaster R., Lanius D., 2021, «Speaking of Fake News: Definitions and Dimensions», in Bernecker S., Flowerree A., Grundmann T. (eds.), *The epistemology of fake news*, Oxford, Oxford University Press, pp. 19-45.

Jodłowiec M., 2023, «Deceptive clickbaits in the relevance-theoretic lens: What makes them similar to punchlines», *Pragmatics*, 33, 3, pp. 418-435.

Johnson L. D., 2024, «Climate change data is based on FRAUD, and scientists around the world are pushing back against the narrative», *Natural News*, https://www.naturalnews.com/2024-03-12-climate-change-data-based-on-fraud.html (last access: 12/10/2024).

Kent S., 2020, «Police in Australia Begin Massive Criminal Investigation into Bushfire Arson», *Breitbart*, https://www.breitbart.com/politics/2020/01/06/police-in-australia-begin-massive-criminal-investigation-into-bushfire-arson/ (last access: 12/12/2024).

Khoo J., 2017, «Code words in political discourse», *Philosophical Topics*, 45, 2, pp. 33-64.

Kingsbury J., McKeown-Green J., 2009, «Definitions: Does Disjunction Mean Dysfunction?», *Journal of Philosophy*, 106, 10, pp. 568-585.

Koechley P., 2012, «Why the Title Matters More Than the Talk», *Medium*, Online: https://blog.upworthy.com/why-the-title-matters-more-than-the-talk-867d08b75c3b (last access: 10/10/2024).

Koehler D. J., 2016, «Can journalistic 'false balance' distort public perception of consensus in expert opinion?», Journal of Experimental Psychology: Applied, 22, 1, pp. 24–38.

Kovach B., Rosenstiel T., 2014, *The Elements of Journalism: What Newspeople Should Know and the Public Should Expect*, New York, Crown Trade.

Kulin J., Johansson Sevä I., Dunlap R. E., 2021, «Nationalist ideology, rightwing populism, and public views about climate change in Europe», *Environmental Politics*, 30, 7, pp. 1111–1134.

Kunda Z., 1987, «Motivated inference: Self-serving generation and evaluation of causal theories», *Journal of Personality and Social Psychology*, 53, 4, pp. 636–647.

Ladyman J., 2013, «Toward a Demarcation of Science from Pseudoscience», in Pigliucci m., Boudry M. (eds.), *Philosophy of Pseudoscience: Reconsidering the Demarcation Problem*, Chicago, University of Chicago Press, pp. 45-60.

Lazer D. M., Baum M. A., Benkler Y., Berinski A. J., Greenhill K. M., Menczer F.,Schudson M., 2018, «The Science of Fake News», *Science*, 359, 6380, pp. 1094-1096.

Leòn J. A., 1997, «The effects of headlines and summaries on news comprehension and recall», *Reading and Writing: An Interdisciplinary Journal*, 9, pp. 85-106.

Lepore J., 2018, *These Truths: A History of the United States*, W. W. Norton & Company: New York.

Levy N., 2017, «The Bad News About Fake News», *Social Epistemology Review and Reply Collective*, 6, 8, pp. 20-36.

Lewandowsky S., 2021, «Climate change disinformation and how to combat it», Annual Review of Public Health, 42, pp. 1–21.

Lewandowsky S., Ecker U., Cook J., 2017, «Beyond Misinformation: Understanding and Coping with the 'Post-Truth' Era», Journal of Applied Research in Memory and Cognition, 6, 4, pp. 353–369.

Lewandowsky S., Gignac G. E., Oberauer K., 2013, «The Role of Conspiracist Ideation and Worldviews in Predicting Rejection of Science», PLOS ONE, 8, 10, e75637, pp. 1-11.

Lewis D., 1969, *Convention: A Philosophical Study*, Cambridge (MA), Wiley-Blackwell.

Lorenz E., 1995, «Climate is what you expect», Prepared for publication by NCAR. Unpublished, pp. 1–33.

Lutzke L., Drummond C., Slovic P., Arvai J., 2019, «Priming critical thinking: Simple interventions limit the influence of fake news about climate change on Facebook», Global Environmental Change, 58, 101964, pp. 1-8.

Lynas M., Houlton B., Perry S., 2021, «Greater than 99% consensus on human-caused climate change in the peer-reviewed scientific literature», Environmental Research Letters, 16, 114005, pp. 1-7.

Mahon J. E., 2016, «The Definition of Lying and Deception», *The Stanford Encyclopedia of Philosophy*, https://plato.stanford.edu/entries/lying-definition/ (last access: 31/07/2024)

Marsili N., 2021, «Retweeting: its linguistic and epistemic value», *Synthese*, 198, pp. 10457–10483.

Marsili N., 2024, «The definition of assertion: Commitment and truth», *Mind & Language*, pp. 1-21, https://doi.org/10.1111/mila.12476.

Mascitti M., 2023, «Dogwhistles and audience design: a new definition», *Manuscrito*, 46, 3, pp. 1-37.

Mazzarella D., 2013, «Optimal relevance as a pragmatic criterion: the role of epistemic vigilance», *UCL Work. Pap. Linguist*, 25, pp. 20–45.

Mazzarella D., 2021, «"I didn't mean to suggest anything like that!": Deniability and context reconstruction», *Mind & Language*, pp. 1-19.

Mazzetti M., 2020, «G.O.P.-Led Senate Panel Details Ties Between 2016 Trump Campaign and Russia», *The New York Times*, https://www.nytimes.com/2020/08/18/us/politics/senate-intelligence-russian-interference-report.html (last access: 31/07/2024)

McGrath M., Frank D., 2023, «Propositions», *The Stanford Encyclopedia of Philosophy*, https://plato.stanford.edu/entries/propositions/ (last access: 09/08/2024).

McIntyre L., 2018, *Post-Truth*, Cambridge, The MIT Press.

McMorris-Santoro E., 2014, «White House: Climate Change Is Real And It's Making Your Life Worse Right Now», *BuzzFeed News*, https://www.buzzfeednews.com/article/evanmcsan/white-house-climate-change-is-real-and-its-making-your-life (last access: 12/12/2024).

Medimorec S., Pennycook G., 2015, «The language of denial: Text analysis reveals differences in language use between climate change proponents and skeptics», Climatic Change, 133, 4, pp. 597–605.

Mercier H., 2020, *Not Born Yesterday: The Science of Who We Trust and What We Believe*. Princeton, Princeton University Press.

Michael R. B., Breaux B. O., 2021, «The relationship between political affiliation and beliefs about sources of "fake news"», *Cognitive Research: Principles and Implications*, 6, pp. 1-15.

Mikkelson D., 2017, «Did Keanu Reeves Say Hollywood Elites Use the 'Blood of Babies' to Get High?», *Snopes*, https://www.snopes.com/fact-check/keanu-reeves-blood-drinking/ (last access: 31/07/2024)

Millar R., Fuglestvedt J., Friedlingstein P., et al., 2017, «Emission budgets and pathways consistent with limiting warming to 1.5 °C», Nature Geoscience, 10, pp. 741–747.

Milman O., 2024, «World's 1.5C climate target 'deader than a doornail', experts say», *The Guardian*, https://www.theguardian.com/environment/2024/nov/18/climate-crisis-world-temperature-target (last access: 12/12/2024).

Mitchell J. F. B., Karoly D., Hegrel G. C., Zwiers F. W., Allen M. R., Marengo J., 2001, «Detection of climate change and attribution of causes», in Houghton J. T., Ding Y., Griggs D. J., Noguer M., van der Linden P. J., Dai X., Maskell K., Johnson C. A. (eds.), Climate Change 2001: The Scientific Basis, Cambridge, Cambridge University Press, pp. 696–738.

Mukerji N. S., 2018, «What is fake news?», *Ergo*, 5, pp. 923–946.

Nelson D. L., McEvoy C. L., Dennis S., 2000, «What is free association and what does it measure?», *Memory & Cognition*, 28, 6, pp. 887–899.

North G. R., 2013, «More evidence for anthropogenic influence on climate change», Proceedings of the National Academy of Sciences of the United States of America, 110, 43, pp. 17169–17170.

Nyhan B., Reifler J., 2010, «When corrections fail: The persistence of political misperceptions», Political Behavior, 32, 2, pp. 303–330.

Origgi G., 2008a, «Trust, authority and epistemic responsibility», *Theoria*, 23, 1, pp. 35–44.

Origgi G., 2008b, «A stance of trust», *Estudios en homenaje a José Luis Guijarro Morales*, pp. 187–200.

Oshikawa, R., Qian, J., & Wang, W. Y. (2020), «A survey on natural language processing for fake news detection», *Proceedings of the 12th Language Resources and Evaluation Conference*, pp. 6086-6093.

Oyserman D., Dawson A., 2020, «Your fake news, our facts: Identity-based motivation shapes what we believe, share, and accept», in Greifeneder R., Jaffe M., Newman E., Schwartz N. (eds.), *The Psychology of Fake News: Accepting, Sharing, and Correcting Misinformation*, London, Routledge, pp. 173-195.

Pagin P., 2021, «Assertion», The Stanford Encyclopedia of Philosophy, https://plato.stanford.edu/entries/assertion/ (last access: 09/08/2024).

Palma B., 2020, «Were 'Nearly 200' People Arrested for Deliberately Starting Australia Bushfires?», *Snopes*, https://www.snopes.com/fact-check/australia-arrested-bushfires/ (last access: 12/12/2024).

Partee B. H., 2021, «Formal Semantics, Lexical Semantics, and Compositionality: The Puzzle of Privative Adjectives», *Philologia*, *7*, 1, pp. 1-24.

Partee B. H., 2021, *Formal Semantics, Lexical Semantics, and Compositionality: The Puzzle of Privative Adjectives*, *Philologia*, 7, 1, pp. 1-24.

Pepp J., 2018, «Truth Serum, Liar Serum, and Some Problems About Saying What You Think is False», In Michaelson E., Stokke A. (eds.), *Lying: Language, Knowledge, Ethics, and Politics*, Oxford, Oxford University Press (2018), pp. 43-64.

Pepp J., Michaelson E., Sterken K. R., 2019, «What's new about fake news?», *Ethics and Social Philosophy,* 16, 2, pp. 67–92.

Pepp J., Michaelson E., Sterken K. R., 2019, «Why we should keep talking about fake news», *Inquiry: An Interdisciplinary Journal of Philosophy*, 65, 4, pp. 471-487.

Piazza T., Croce M., 2019, «Epistemologia delle fake news», *Sistemi Intelligenti*, 31, 3, pp. 433–461.

Pinker S., Nowak A. M., Lee J. J., 2008, «The logic of indirect speech», *Proceedings of the National Academy of Sciences of the United States of America*, 105, pp. 833-838.

Pollaj J. B., 2023, «San Francisco Braces for First Snow Since 1976 as More Winter Storms Arrive», *Breitbart*, https://www.breitbart.com/local/2023/02/22/san-francisco-braces-for-first-snow-since-1976-as-more-winter-storms-arrive/ (last access: 12/12/2024).

Potthast M., Kiesel J., Reinartz K., Bevendorff J., Stein B., 2018, «A Stylometric Inquiry into Hyperpartisan and Fake News», Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, 1, pp. 231-240.

Powell J., 2017, «Scientists reach 100% consensus on anthropogenic global warming», Bulletin of Science, Technology & Society, 37, 4, pp. 183–184.

Pritchard D., 2021, «Good News, Bad News, Fake News», in Bernecker S., Flowerree A. Grundmann T. (eds.), *The epistemology of fake news*, Oxford, Oxford University Press, pp. 46-67.

Puglisi R., Snyder M. J., 2015, «Empirical studies of Media Bias», in Anderson P. S., Waldfogel J., Strömberg D. (eds.), *Handbook of Media Economics*, North-Holland, Elsevier, pp. 647-667.

Pustejovsky J., 1995, *The generative lexicon*, Cambridge (MA), The MIT Press.

Putnam H., 1975, *Mind, Language and Reality: Philosophical Papers*, Cambridge University Press: New York.

Ranney M. A., Clark D., 2016, «Climate Change Conceptual Change: Scientific Information Can Transform Attitudes», Topics in Cognitive Science, 8, 1, pp. 49–75.

Readfearn G., 2016, «Revealed: Most Popular Climate Story on Social Media Told Half a Million People the Science Was a Hoax», DeSmog, https://www.desmog.com/2016/11/29/revealed-most-popular-climate-story-social-media-told-half-million-people-science-was-hoax/ (last access: 10/12/2024).

Reed M., 2016, «'This loopy idea': An analysis of UKIP's social media discourse in relation to rurality and climate change», Space and Polity, 20, 2, pp. 226–241.

Reuters Fact Check, « Fact Check: No evidence that people aged 25-44 experienced an 84% increase in excess mortality due to COVID vaccine rollout», *Reuters*, Online: https://www.reuters.com/article/fact-check/fact-check-no-evidence-that-people-aged-25-44-experienced-an-84-increase-in-ex-idUSL2N2VS1BI/ (last access: 10/10/2024).

Rini R., 2017, «Fake news and partisan epistemology», *Kennedy Institute of Ethics Journal*, 27, 2S, pp. E-43–E-64.

Rogelj J., et al., 2016, «Paris Agreement climate proposals need a boost to keep warming well below 2°C», *Nature*, 534, 7609, pp. 631–639.

Rogelj, J., den Elzen M., Höhne N., Fransen T., Fekete H., Winkler H., Schaeffer R., Sha F., Riahi K., Meinshausen M., 2016, «Paris Agreement climate proposals need a boost to keep warming well below 2 °C», *Nature*, 534, 7609, pp. 631-639.

Rubin V. L., Chen Y., Conroy N. K., 2015, «Deception detection for news: Three types of fakes», *Proceedings of the Association for Information Science and Technology*, 52, 1, pp. 1-4.

Ryan S., 2021, «Fake News, Epistemic Coverage and Trust», *The Political Quarterly*, 92, pp. 606-612.

Saul J., 2018, «Dogwhistles, Political Manipulation, and Philosophy of Language», in Harris D., Fogal D., Moss M. (eds.), *New Work on Speech Acts*, Oxford, Oxford University Press, pp. 360-383.

Scarantino A., Piccinini G., 2010, «Information without Truth», *Metaphilosophy*, 41, 3, pp. 313-330.

Schweikard P. D., Hans Bernhard Schmid B. H., 2020, «Collective Intentionality», *The Stanford Encyclopedia of Philosophy*, https://plato.stanford.edu/cgi-bin/encyclopedia/archinfo.cgi?entry=collective-intentionality (last access: 09/08/2024).

Scott K., 2021, «You won't believe what's in this paper! Clickbait, relevance and the curiosity gap», *Journal of Pragmatics*, 175, 53-66.

Scott K., 2023, «"Deceptive" clickbait headlines: Relevance, intentions, and lies», *Journal of Pragmatics*, 218, pp. 71-82.

Searle J. R., 1969/1992, *Atti linguistici. Saggio di filosofia del linguaggio*, Torino, Bollati Boringhieri.

Searle R. J., 1990, «Collective Intentions and Actions», In Cohen R. P., Morgan J., Pollack M. (eds.), *Intentions in Communication*, Cambridge, MIT Press, pp. 401-415.

Searle R. J., 1995, *The Construction of Social Reality*, New York, The Free Press.

Searle R. J., 2005 «What is an institution?», *Journal of Institutional Economics*, 1, 1, pp. 1-22.

Searle R. J., 2010, Making the social world: The Structure of Human Civilization, Oxford, Oxford University Press.

Sharman A., 2013, «Mapping the climate sceptical blogosphere», GRI Working Papers, 124, https://ideas.repec.org/p/lsg/lsgwps/wp124.html.

Shu K., Sliva A., Wang S., Tang J., Liu H., 2017, «Fake news detection on social media: A data mining perspective», *ACM SIGKDD Explorations Newsletter*, 19, 1, pp. 22-36.

Silverman C., 2016, «I Helped Popularize The Term "Fake News" And Now I Cringe Every Time I Hear It», *BuzzFeed News*,

https://www.buzzfeednews.com/article/craigsilverman/i-helped-popularize-the-term-fake-news-and-now-i-cringe (last access: 31/07/2024)

Silverman C., 2016, «This Analysis Shows How Viral Fake Election News Stories Outperformed Real News On Facebook», *BuzzFeed News*, https://www.buzzfeednews.com/article/craigsilverman/viral-fake-election-news-outperformed-real-news-on-facebook (last access: 10/10/2024).

Silverman C., Lawrence A., 2016, «How Teens In The Balkans Are Duping Trump Supporters With Fake News», https://www.buzzfeednews.com/article/craigsilverman/how-macedonia-became-a-global-hub-for-pro-trump-misinfo (last access: 31/07/2024).

Simion M., 2023, «Knowledge and Disinformation», *Episteme*, pp. 1-12, doi:10.1017/epi.2023.25.

Smith B., Searle R. J., 2003, «The Construction of Social Reality: An Exchange», *American Journal of Economics and Sociology*, 62, 2, pp. 285-309.

Smith E. J., 1999, «Leadlines may be better than traditional headlines», *Newspaper Research Journal*, 20, pp. 55-64.

Smith E. K., Mayer A., 2019, «Anomalous Anglophones? Contours of free market ideology, political polarization, and climate change attitudes in English-speaking countries, Western European and post-Communist states», Climatic Change, 152, pp. 17–34.

Soll J., 2016, «The Long and Brutal History of Fake News», *Politico*, https://www.politico.com/magazine/story/2016/12/fake-news-history-long-violent-214535/ (last access: 31/07/2024)

Sperber D., Clement F., Heintz C., Mascaro O., Mercier H., Origgi G., Wilson D., 2010, «Epistemic Vigilance», *Mind and Language*, 25, 4, pp. 359–393.

Sperber D., Wilson D., 1987, «Presumptions of relevance», *Behavioral and Brain Science*, 10, pp. 736-754.

Sperber D., Wilson D., 1995, *Relevance theory: Communication and Cognition (2nd ed.)*, Hoboken, Blackwell Publishing.

Stalnaker R., 2002, «Common ground», vol. 25: *Linguistics and Philosophy*, 5, 6, pp. 701–721.

Stokke A., 2013, «Lying, Deceiving, and Misleading», *Philosophy Compass*, 8, 4, pp. 348-359.

Stone D. A., Knutti R., 2010, «Weather and climate», in Fung F., Lopez A., New M. (eds.), Modelling the Impact of Climate Change on Water Resources, Chichester, John Wiley & Sons, pp. 4-33.

Sullivan M., 2013, «Repairing the Credibility Cracks», *The New York Times*, https://www.nytimes.com/2013/05/05/public-editor/repairing-the-credibility-cracks-after-jayson-blair.html (last access: 10/10/2024).

Talisse R., 2018, «There's No Such Thing as Fake News (And That's Bad News)», *3AM Magazine*, https://www.3ammagazine.com/3am/theres-no-such-thing-as-fake-news-and-thats-bad-news/ (last access: 31/07/2024).

Tandoc E. C., Zheng W. L., Ling R., 2018, «Defining "Fake News": A Typology of Scholarly Definitions», *Digital Journalism*, 6, pp. 137-153.

Thayer E. R., 1989, The Biopsychology of Mood and Arousal, Oxford, Oxford University Press.

Tong C., Gill H., Li J., Valenzuela S., Rojas H., 2020, «"Fake News Is Anything They Say!" — Conceptualization and Weaponization of Fake News among the American Public», *Mass Communication and Society*, 23, 5, pp. 755-778.

van der Linden S., Panagopoulos C., Roozenbeek J., 2020), «You are fake news: political bias in perceptions of fake news» *Media, Culture & Society*, 42, 3, pp. 460-470.

van der Linden S., Roozenbeek J., 2020, «Psychological Inoculation Against Fake News», in Greifeneder R., Jaffe M., Newman E., Schwartz N. (eds.), *The Psychology of Fake News: Accepting, Sharing, and Correcting Misinformation*, London, Routledge, pp. 147-169.

van Eck C. W., Mulder B. C., van der Linden S., 2021, «Echo Chamber Effects in the Climate Change Blogosphere», Environmental Communication, 15, 2, pp. 145–152.

Vosoughi S., Roy D., Aral S., 2018, «The spread of true and false news online», *Science*, 359, pp. 1146-1151.

Vosoughi S., Roy D., Aral S., 2018, «The spread of true and false news online», Science, 359, pp. 1146–1151.

Walter S., Brüggemann M., Engesser S., 2017, «Echo Chambers of Denial: Explaining User Comments on Climate Change», Environmental Communication, 12, pp. 1–14.

Wang L., Ramachandran A., Chaintreau A., 2021, «Measuring Click and Share Dynamics on Social Media: A Reproducible and Validated Approach», *Proceedings of the International AAAI Conference on Web and Social Media*, 10, 2, pp. 108-113.

Waterson J., 2020, «'They're building a massive lasagne': man behind WhatsApp virus spoof revealed», *The Guardian*, https://www.theguardian.com/technology/2020/mar/21/a-lasagne-the-size-of-wembley-whatsapp-joke-swallowed-whole (last access: 31/07/2024).

Watson D., Wiese D., Vaidya J., Tellegen A., 1999, «The two general activation systems of affect: Structural findings, evolutionary considerations, and psychobiological evidence», Journal of Personality and Social Psychology, 76, 5, pp. 820–838.

Werndl C., 2015, «On defining climate and climate change», The British Journal for the Philosophy of Science, 67, 2, pp. 337-364

Wilson D., Sperber D., 2004, «Relevance Theory», in Horn L. R., Ward G. (eds.), *The Handbook of Pragmatics*, Oxford, Blackwell, 607-632.

Witten K., 2014, «Dogwhistle politics: the new pitch of an old narrative», Unpublished manuscript, downloaded from academia.edu.

Wobbrock J. O., Hattatoglu L., Hsu A. K., Burger M. A., Magee M. J., 2021, «The Goldilocks zone: young adults' credibility perceptions of online news articles based on visual appearance», *New Review of Hypermedia and Multimedia*, 27, 1-2, pp. 51-96.

Wood T., Porter E., 2019, «The elusive backfire effect: Mass attitudes' steadfast factual adherence», Political Behavior, 41, 1, pp. 135–163.

Wreen M., 1980, «Counterfeit Coins and Forged Paintings: Caveat Emptor», *Analysis* , 40, 3, pp. 146-151.

Wreen M., 2002, «Forgery», *Canadian Journal of Philosophy*, 32, 2, pp. 143-166.

Wright C., 1992, *Truth and Objectivity*, Cambridge, Harvard University Press.

Xu Z., Laffidy M., Ellis L., 2022, «Clickbait for climate change: Comparing emotions in headlines and full-texts and their engagement», Information, Communication & Society, 26, pp. 1915-1932.

Zhou X., Zafarani R., 2020, «A survey of fake news: Fundamental theories, detection methods, and opportunities», *ACM Computing Surveys (CSUR)*, 53, 5, pp. 1-40.



The PhD Fellowship was co-financed with funds from the National Operational Programme for Research and Innovation 2014-2020, by FSE REACT-EU funds Action IV.4 “PhD Courses and Research contracts on innovation topics” and Action IV.5 “PhD Courses on Green related topics”