

Epistemic Dependence on Official Experts

Master's Thesis submitted for the

Research Master's Philosophy

Utrecht University

May 29, 2018

Maria Giulia Napolitano

Student ID: 5970806

Supervisor: Prof. Dr. Daniel Cohnitz

Second Reader: Dr. Janneke van Lith

Third Reader: Prof. Dr. Mauro Bonazzi



Universiteit Utrecht

Abstract

In our epistemic lives, we need to depend on experts to a great extent in order to form our opinions in highly specialised domains. But how exactly are experts' opinions to be taken into account by laypeople? And how can laypeople identify reliable experts? In this paper I propose an account of epistemic dependence on experts aimed at answering these questions by focussing on the way in which we encounter expertise in our everyday lives, through the notion of *official experts*. I employ the case of vaccination scepticism as a means of illustration of a controversial epistemic situation, in order to investigate the conditions under which it is reasonable for laypeople to mistrust official experts, and propose possible ways in which laypeople can assess expertise without having to engage with the content of the claims made by experts in their domains. Through a discussion of epistemic dependence on experts, my aim is to show on epistemological grounds alone why it is irrational for laypeople to disagree with the majority of official experts in cases such as the vaccination controversy.

Word Count: 26 123

Table of Contents

Introduction.....	1
1. Epistemic authorities: Linda Zagzebski and Thomas Grundmann on preemption.....	6
2. Official experts as epistemic authorities.....	12
3. Dependence on the community of official experts for coverage	20
4. How can laypeople assess conflicting authorities' reports?.....	28
5. Experts disagreement: going by the numbers.....	36
6. Doubts about official experts' epistemic coverage: when is a rumour reliable?....	44
Conclusions.....	52
Bibliography.....	55

Introduction

We live in societies that are based on a thorough division of cognitive labour: we all know many things, but only have direct evidence of a fragment of them. We rely on the discoveries of the past, and we trust people who have greater competence in certain fields— the experts. There is too much to be known, too much is yet to be investigated, and we are aware of having limited cognitive and material resources to do all of this on our own. That's why we trust some members of our communities to do this for us, each in their own specialised field of expertise, and we rely on them to form our opinions. In other words, we appeal to *epistemic authorities* in order to form our beliefs about the world in those fields in which we have reasons to believe them to be in a better epistemic position to arrive at the truth than we are. Hence, we trust doctors to cure us from diseases, engineers to build our houses, journalists to report the news, scientists in every field to produce new knowledge and to apply it in order to make our lives better. Call this phenomenon *epistemic dependence on experts*. However, sometimes the reliance on experts does not go as smoothly as this: sometimes laypeople are sceptical of what the experts say, and don't know what to believe and whom to trust. This phenomenon is possibly amplified in our times by internet and social media, that make the spread and retrieval of all kinds information, including information that falls within highly specialised domains of expertise, easy, quick and available to everyone. Every internet user can consult different putative experts' opinions online, and discuss specialised topics with other users, while the filtering of this massive load of information is only left to their personal judgement. Under these circumstances, it can be extremely difficult for laypeople to decide what to believe for domains in which expertise is required to form a reliable opinion.

As an instance of this, I want to consider as an example a very actual and worrisome phenomenon: the case of vaccination scepticism. Against the scientific consensus that vaccines are safe and effective, many parents decide not to vaccinate their children fearing harmful side effects, and sometimes doubting their efficacy. According to a study in 2016 Europe was the region with the lowest confidence in vaccination safety (Larson et al. 2016). Immunisation rates are still not high enough to guarantee herd immunity, necessary to prevent disease outbreaks.¹ With the aid of social media, anti-vaccination

¹ <http://www.euro.who.int/en/media-centre/sections/press-releases/2018/europe-observes-a-4-fold-increase-in-measles-cases-in-2017-compared-to-previous-year>

movements share studies about vaccination safety, efficacy, and practical information about vaccination legislation and how to avoid mandatory vaccination, where present.² Laypeople are presented with loads of information on both sides of the dispute, and it can be hard to decide whom to trust and what to believe. For the purpose of this paper, it is not necessary to go into the factual details of the vaccine controversy, and the case will only be considered as a means of illustration of a controversial case of dependence on experts. The focus is rather the phenomenon of vaccination scepticism *as such*, i.e., the fact that some people, who are not experts in medicine or biology, believe to have reason to mistrust the vast majority of the official experts in those fields, and believe, against the experts' consensus, that vaccines are harmful or not efficacious.³ The case of vaccination scepticism will serve as an example to guide our analysis of epistemic dependence on experts. My aim in this paper is to show, on *epistemological grounds* alone, how laypeople should go about forming their opinion on topics such as the safety and efficacy of vaccinations. I will *not* consider the question whether vaccinations are actually safe and efficacious, but only *what laypeople have reason to believe* about vaccines— and in fact about *any other claim* made by official experts. I shall do this by analysing the epistemology of the dependence on experts.

Three sets of issues are central to the topic of epistemic dependence on experts. First, as laypeople, we need to be able to *identify* the experts, and this can be a challenging task, given that the reason to trust experts in the first place is that we lack knowledge in that field. As laypeople we are rarely in a position to assess experts on the basis of the *claims* they make. Being an ignoramus in the field, I have to believe my mechanic when he tells me that my car is making that grinding noise because the braking pads need to be replaced. For all I know, he could just as well have said that the reason is the engine belt, and that would not have given me any indication of his expertise. In order to identify

² For Europe: <https://www.efvv.eu/>

³ Surely, the problem of vaccination cannot be simply *reduced* to a problem with the reliance on experts: vaccination-resistance is a complex phenomenon, that can be traced back to a multiplicity of factors, including conflicts of values (such as liberty and purity: Amin et al. 2017), the resistance to a policy choice (herd immunisation over individual risks), and belief in alternative medicine. It is moreover questionable whether the decrease in the vaccination rates can be attributed to a significant degree to vaccination-resistance, or if it is rather the effect of economic and social factors preventing people from accessing health care. However, one thing remains: despite the scientific consensus, many people nowadays are sceptical of vaccines' efficacy and safety and are hesitant to get their children vaccinated.

experts, it seems that we have to rely on indirect indicators, such as the reputation of my mechanic, or their ‘accessible’ predictions (e.g., if the car broke down two days after he ‘fixed’ it, I would probably question my mechanic’s expertise). Secondly, after having identified the experts, there is the issue of *how it is rational to respond* to their claims, i.e., what weight should be given to the expert’s opinion when we form our opinion on the topic: is the fact that my mechanic thinks that the noise is due to the braking pads an additional reason to believe the braking pads are the cause of the noise, to be added to other reasons I may have (e.g., that I read on the internet that a grinding noise is usually due to overused braking pads)? Or should I just blindly follow the expert, and be just as confident as my mechanic is that the problem is in the braking pads, just because ‘he says so’? In other words, is the expert’s opinion to be treated as a piece of additional evidence to be weighted against the other relevant evidence we may have, or is it a preemptive reason to believe what the expert claims, that screens off all of our other relevant evidence? Finally, closely related to this last point, there is the issue of *under what circumstances it is rational not to rely* on the expert’s opinion on a claim that falls within their domain of expertise. For instance, I may be sceptical of what my mechanic says if he was clearly drunk while checking my car, or if another mechanic told me that the noise was not due to the braking pads, but rather to the engine. But why is this so? And what should I believe in these cases? In this paper I want to provide an analysis of the phenomenon of epistemic dependence on experts, by addressing these three issues in a unified way, and apply this analysis to the case study of vaccination scepticism.

In the literature, the problem of identifying experts has been treated separately from the question of how we should incorporate the opinion of experts in our judgments. Alvin Goldman has been the first to systematically address the problem of how it is possible for laypeople to identify experts in a field (1999, ch. 8; 2001)⁴. According to him, non-experts can assess expertise by appealing to two main sources: the first, which is defended in *Knowledge in a Social World*, is the possibility of verification of a putative expert’s claims, both by directly verifying the expert’s claim, or by assessing the expert’s predictions (1999: 268-271); the second source for identifying expertise (in a comparative way, when a layperson has to decide between two putative experts) consists of indirect indicators of expertise, which he identifies with (i) argumentative superiority, (ii)

⁴ The problem was already formulated by Hardwig (1985; 1991). He does not, however, point at possible ways in which laypeople could reliably identify experts.

agreement from other experts, (iii) appraisal by ‘meta-experts’ (e.g., credentials), (iv) evidence of interests and biases, (v) evidence of past track-records (2001: 93). Goldman does not engage with the issue of how the non-expert should respond to experts, and how her own opinion on the subject matter should or should not affect her final judgement and the identification of the expert. Goldman relies on an idealised case where the non-expert ‘is not in a position to evaluate the target experts by using his own opinion; at least he does not think he is in such a position’ (Goldman 2001: 90), thus leaving open the possibility that a layperson *may* use her own opinions on a subject matter to identify an expert, even when she knows that she is not an expert in the field. For instance, being confident of my opinion that the noise my car is making is due to the engine belt, even though I know I am not an expert when it comes to cars, I may question my mechanic’s expertise when he tells me that the noise is due to the braking pads. In order to assess whether it is rational for a layperson to let her object-level beliefs in a domain interfere with her assessment of an expert’s expertise in the same domain, it is helpful to look at how laypeople should respond to the testimony of experts.

The debate on this topic is dominated by two conflicting views. Some epistemologists believe that experts’ opinion should be considered as additional evidence, to be added to and weighted against other relevant evidence we may have on the topic— this view is known as the *total evidence view*⁵. Opposed to the total evidence view is the *preemption view*, according to which experts’ opinions should be endorsed on authority, upon recognising that they are in a better epistemic position than we are. According to the preemption view, the opinion of an epistemic authority screens off all the additional evidence we may have to decide on a certain claim and forces us to adopt the expert’s belief just on the basis that it is what the authority believes⁶. Different stances on the question how experts’ opinions should be taken into account lead to different positions on the third core issue of epistemic dependence on experts, namely the circumstances under which it is rational *not* to defer to the experts. The research on these topics, however, is not connected to the problem of identifying experts. Philosophers from both sides of the

⁵ Supporters of different versions of the total evidence view are Jäger (2016); Lackey (2018); Dormandy (forthcoming).

⁶ This is the view defended by Zagzebski (2012; 2013). Other accounts of preemption are defended by Keren (2007; 2014a; 2014b) Grundmann (manuscript). The preemption view resembles what in formal epistemology has been defended as ‘deference model’. See Elga (2007); Joyce (2007).

divide rely on a (more or less spelled out) definition of authorities and experts as people who are better suited to arrive at knowledge in a certain field. They assume that we know who the experts are, and ask how it is rational to incorporate their opinion when judging about certain matters that fall within their domain of expertise. Nevertheless, the rational way to respond to experts crucially depends on how experts and authorities are defined, and on whether we have reasons to believe we are reliable at identifying them.

In the first part of this paper I argue that, once the core concepts of experts and authorities are appropriately defined, and the ‘mechanism’ of preemption is explained in familiar epistemological terms, the preemption view clearly proves to be the correct one. Moreover, my goal is to provide an account of epistemic dependence on experts that is grounded in the way we actually encounter expertise in our lives, in the form of *official experts*. I shall thus analyse the notion of official experts and the reasons we have to treat them as epistemic authorities. Through an analysis of official experts, my aim is to bridge the gap between the two discussions regarding the identification of experts and how we should respond to experts’ opinions. If my account is correct, the social indicator of ‘official expertise’ provides the layperson with *prima facie* reasons to treat official experts as authoritative. The account of epistemic dependence on experts that I outline in the first part of the paper will serve as a starting point to investigate, in the second part, possible reasons for laypeople to question the authoritativeness of official experts’ claims and silence, and possible ways in which these doubts can be addressed by laypeople without them having to acquire expertise in the field. The result of this analysis is an account of epistemic dependence on experts that, taken as a whole, should be sufficient *for laypeople* to decide about the rationality of vaccination scepticism and similar controversies on epistemological grounds alone.

The paper is structured as follows: in the first section I discuss Zagzebski’s pioneering account of epistemic authority (2012; 2013), and the way in which Thomas Grundmann (manuscript) has improved on it by explaining preemption as an instance of undercutting defeat and by identifying the defining characteristics of epistemic authority. In the second section, I introduce the notion of official experts, and I develop my own account of expertise and authority, taking the cue from Grundmann’s general ideal, but twisting some of his notions to integrate his account with an analysis of official expertise. I then argue that we have *prima facie* reasons to treat official experts as epistemic authorities for every question within their domain of expertise. In the third

section, I present another dimension of the epistemic reliance on experts: the reliance on the community of experts for coverage. Laypeople rely on experts to reliably investigate new claims and communicate about discoveries in their domain of expertise, and this implies that experts' silence is epistemically significant for laypeople, just as experts' testimony is. In the fourth section, I discuss possible reasons to mistrust official experts, and the sources of evidence available for laypeople to assess expertise when experts disagree. I highlight the limited applicability of some of the methods proposed by Goldman (2001), namely argumentative superiority, evidence of biases and interests, and evidence of experts' past track record. In the fifth section, I discuss the method of going by the numbers as the most reliable source for laypeople for assessing competing authorities. I discuss the criterion of independence (Goldman 2001) in light of some recent attacks to it (Coady 2006; Lackey 2013). Finally, in the last section, I focus on possible reasons for laypeople to be sceptical of the epistemic coverage provided by experts. In particular, I consider whether rumour can be evidence for the truth of a domain specific claim. The case of vaccination scepticism will be employed as a recurring example, and will be in the background in every section, as a motivation for this analysis, and as a thread to guide us through its various steps.

1. Epistemic authorities: Linda Zagzebski and Thomas Grundmann on preemption

In the case of vaccinations, many parents find themselves in the following epistemic situation: they are presented with the opinion of an expert, such as their family doctor, and they need to form their own opinion on the basis of this. According to the *total evidence view*, the doctor's opinion counts as one source of evidence, to be weighted against other evidence they may have to make up their minds on the topic, such as information gathered on the internet, and the opinions of their friends. On the other hand, according to the *preemption view*, parents should just follow the doctor's advice on the basis of her *authority*, and all the other evidence they may have in favour or against her opinion is normatively screened off. In order to articulate my account of epistemic dependence on expert, I will start by discussing— and ultimately defending— the preemption view of epistemic authority.

In her book “Epistemic Authority: a Theory of Trust, Authority and Autonomy in Belief”, Zagzebski defended an influential account of epistemic authority, modelled on

Joseph Raz's account of political authority (Raz 1988), according to which it is sometimes rational for a person to believe what another person says preemptively, just because the other person believes it. Following Raz, Zagzebski defines an epistemic authority as an agent for whom four conditions hold: the *Content Independence Thesis*, the *Preemption Thesis*, the *Dependency Thesis* and the *Normal Justification Thesis*. According to the Content Independence Thesis, an authority is such if there is no direct connection between the reason expressed by the authority's utterance, and the belief for which it is a reason. In other words, for an epistemic authority it is the case that "if the epistemic authority had believed a different proposition, the subject would have had reason to believe the other proposition instead" (Zagzebski 2013: 297). If the authority asserts 'p', this gives you a reason to believe p, but if she asserted $\neg p$, then you should have believed $\neg p$: the normative force that the authority possesses does not depend on the content of her utterances. The second condition that defines epistemic authority is the Preemption Thesis:

Pre-emption Thesis for Epistemic Authority: The fact that the authority has a belief p is a reason for me to believe p that replaces my other reasons relevant to believing p and is not simply added to them. (Zagzebski 2013: 298)

An epistemic authority saying 'p' is not just evidence in favour of p: an authority saying p is a reason to believe p that screens off the other reasons one might have for believing p or $\neg p$.

The other two conditions of epistemic authority are aimed at establishing the rationality of preemption. In fact, according to the Dependency Thesis, the beliefs that a person takes to be authoritative should be formed in a way that I take worthy of emulation. Moreover, according to the Justification Thesis:

The authority of another person's belief for me is justified by my conscientious judgment that I am more likely to form a true belief and avoid a false belief if I believe what the authority believes than if I try to figure out what to believe myself. (Zagzebski 2013: 299)

While the Content Independence Thesis and the Preemption Thesis illustrate how we should respond to an epistemic authority's claim, the Dependency Thesis and the Justification Thesis are meant to explain why it is rational to do so: If I believe that a certain person has formed a belief in a way that is worthy of emulation, and I believe that I am more likely to be correct about a certain thing if I defer to what this person

believes than if I try to make up my own mind, then this person is for me an *epistemic authority*, and I should adopt her beliefs preemptively, independently of their content.

Zagzebski's account of epistemic authority asserts the possibility of rationally believing what someone else believes, instead of trying to decide for ourselves what to believe about a certain question. However, her point does not seem to go much beyond the observation that, if we want to get things right, and we believe someone to be more likely than us to get things right, we should follow them and believe what they believe. This much seems trivially true. It is much less clear why we should *preempt* our own reasons for believing what the authority believes, instead of adding them to the reasons we may already possess (see Jäger 2016: 177; Lackey 2018: 238; Dormandy forthcoming)⁷. If I already believe *p*, and the authority confirms my belief, it is not clear why I should not base my belief on all the reasons I had, plus the authority's agreement. The principle of preemption in Zagzebski's view looks rather mysterious and is not reduced to other epistemic principles. The second major problem with Zagzebski's account regards the notion of epistemic authority. Zagzebski's definition of epistemic authorities in terms of the normative power to preempt one own's reasons for belief is, at best, incomplete: it remains silent on the characteristics that make epistemic authorities authoritative, and does not explain which features of epistemic authorities make it rational to preempt one own's reasons in favour of the authority's. Nor does it provide an account of how to identify reliable authorities, which is a crucial issue if one aims at improving one own's epistemic position by deferring to a putative authority.

A refined version of the preemption view of epistemic authority has been defended by Thomas Grundmann (manuscript). Grundmann improves on Zagzebski by providing an account of authority that covers both issues mentioned above: first, instead of defining an epistemic authority by her capacity to give preemptive reasons, he identifies characteristics that are ascribed to epistemic authorities and show that in virtue of these characteristics authorities give preemptive reasons; secondly, he explains preemption as an instance of a more general epistemic principle, through the notion of undercutting defeat.

Grundmann links the notion of epistemic authority to the notions of expertise and epistemic superiority. According to him:

⁷ For a reply to this criticism see Zagzebski (2016: 190).

EA: A is an epistemic authority for S with respect to domain D iff S has justification to believe that (i)A is an expert about D; (ii)A is an epistemic superior to S with respect to D.

Grundmann (manuscript: 8,9).

Following Goldman (2001: 91, 92), Grundmann considers expertise to consist of two components: (i) the possession by the expert of a substantive body of evidence and (ii) methods which reliably lead them to form true beliefs in a certain domain. While expertise is an objective measure, epistemic superiority is relational. An epistemic superior is someone who we consider more likely than us to form correct beliefs in a certain domain. Like expertise, superiority is a matter of both possession of evidence and evaluation. An epistemic authority has access to a superior body of evidence and is able to assess the evidence better than a layperson. The authority's superiority at identifying and gathering evidence makes it such that if I have identified an epistemic authority, I have *prima facie* reasons to believe that the authority has considered all the relevant evidence I possess. By defining epistemic authorities in terms of expertise and superiority together Grundmann manages to avoid the problems of considering incompetent superiors as authoritative (because the expertise condition is not met), and of experts having to consider other experts as authorities (because the superiority condition is not met). As regards the domain (D), Grundmann defines it as "the set of propositions that are esoteric for S [the subject] and exoteric for A [the authority]" (manuscript: 12), where an *esoteric* proposition is understood as a proposition that a layperson is less reliable at assessing than the authority due to the authority's more reliable arsenal of methods, and an *exoteric* one as a proposition which can be assessed in an equally reliable way by the authority and the layperson (manuscript: 11; Cfr. Goldman 2001: 94). In other words, a domain of expertise is a set of facts which are best accessible via specific methods (Grundmann, manuscript: 12).

Grundmann considers **EA** to be an account of *grounded* authority, as opposed to Zagzebski's *role* authority: instead of defining authorities by their capacity to give preemptive reasons, Grundmann takes the characteristics identified in EA to explain why epistemic authorities provide us with preemptive reasons. How is this supposed to work? Grundmann identifies preemption with an instance of undercutting defeat. An undercutting defeater for a subject S's belief that p is a *prima facie* reason for S to believe that "relying on E alone in assessing p would not raise the likelihood of S being right about p over the likelihood of her being right based on pure guessing" (Grundmann

manuscript: 15; see also: Pollock 1974; Bergmann 2005). An undercutting defeater is a reason to believe that the evidence is not connected to the truth of p in the right way, and makes it therefore irrational for S to base her belief on that evidence. According to him, recognising someone as an epistemic authority provides an undercutting defeater for holding any doxastic attitude that deviates from the authority's: recognising someone as an epistemic authority *prima facie* provides reasons to believe that this person is both more competent than you at assigning correct credences in a domain, and has considered more evidence than you, including all of your relevant evidence on the subject. Simply put, having a degree of belief in a proposition p that is different from a recognised authority A , is evidence of being mistaken about p . The ways in which one could be mistaken are many: given the authority's superior competence at evaluating the evidence, learning about the authority's credence could indicate that the evidence must have been incorrectly assessed; or it could be that the authority possesses further evidence of which one is unaware; or it could be evidence that one is double-counting evidence by aggregating one own's reasons with the authority's (Grundmann manuscript: 22). As a subject I need not know what exactly went wrong in my evaluation, it is enough to know that an epistemic authority ends up with a different credence than I do, to realise that I must be mistaken, and that it would be irrational for me to keep relying on my body of evidence to decide about p . Grundmann's account allows for exceptions to full preemption in cases where the layperson has reasons to believe that the authority has not considered a certain piece of evidence, or if the layperson possesses evidence that does not fall within an expert's domain of expertise, or if the undermining of evidence is not full but partial, due to the weak reason one might have for considering someone to be an authority (Grundmann manuscript, 24).

Grundmann's account of the preemptive force of epistemic authorities in terms of undercutting defeat improves on Zagzebski's insofar as preemption is explained as an instance of a general epistemological principle, and the normative capacity of epistemic authorities to provide preemptive reasons is traced back to the defining features of authorities (expertise and superiority). However, just like Zagzebski, Grundmann leaves aside the question of how we should identify authorities, and whether we have reasons to believe that we are reliable at doing so. In Zagzebski's proposal this is particularly problematic, since she believes that it is sometimes possible in some cases for object-level claims made by the putative authorities to count as defeaters for believing them to be authorities, as it happens with authorities' outrageous claims (Zagzebski 2013: 302): if

someone who we consider an authority makes an outrageous claim, this can count as a reason to believe that she is not an authority, as long as we are more convinced of the falsity of the authority's claim than we are of the authority's expertise. But if our identification of authorities is affected by our object-level claims, then we have no reasons to believe that our epistemic position is going to improve by relying on the authorities, unless we have good reasons to believe we are reliable at identifying them, i.e., that our object-level beliefs are not too far off track. But how can we know this, if the very reason to rely on epistemic authorities is that as laypeople we are not very reliable in highly specialised cognitive domains? Zagzebski's account of epistemic authority makes preemption a matter of consistency: if we believe that someone is more likely than us to get things right, then we should believe what they believe, but if we are more sure of our object-level beliefs than of the fact that someone is more likely than us to get things right than it is rational *not* to preempt, and to abandon instead the belief that someone is an epistemic authority (2013: 302, 303; see also Elga 2007: 483).

Differently from Zagzebski, Grundmann seems to believe that object-level propositions that fall within an authority's domain of expertise cannot be rationally used to evaluate the authority's expertise. This emerges from his response to the objection regarding authorities' outrageous beliefs. Grundmann argues that it is never rational for a layperson to dismiss an authority's domain relative belief on the basis of her own evidence within the domain: if your doctor tells you to take 4000 pills a day to cure the measles, you cannot dismiss this claim on the basis that it is outrageous (this example was formulated in Zagzebski 2012:116). The reason why you cannot do it, is that knowing that your authority is a superior with regards to both evidence and competence, and is an expert, your own evidence becomes rationally unusable to determine whether it is right to take 4000 pills to cure the measles. Grundmann holds that you could only refrain from deferring to your authority if you had evidence which fell outside of the domain: e.g., if it was physically impossible to take such a high amount of pills in a day, or if you knew that other doctors never made such a prescription for the same symptoms. These reasons are not domain relative and can therefore be used to decide about the authority's claim (Grundmann, manuscript: 35). However, it is not clear why the outrageous claim could not be used to revise *one's belief about the authority's expertise and superiority*. Why is it the case that my doctor telling me to take 4000 pills could not be taken as an indication that she is probably not an expert? The reasons for believing this in Grundmann's account is that we should possess good *independent* reasons (manuscript: 35) to believe that someone is

an epistemic authority for us. But how can we ever be sure of having identified a reliable expert? What are good reasons to believe someone to be an expert, which are independent of the content of their claims?

Both Zagzebski and Grundmann set aside the question of how to identify reliable authorities, and focus on how it is rational to respond to an authority, once we have identified one. However, this question becomes crucial if we are concerned with the question of how laypeople can improve their epistemic position in highly specialised domains by relying on experts, such as the vaccination controversy. If, like Zagzebski, we allow object level claims to orient our identification of experts, we are not likely to significantly improve our position as laypeople: if, for instance, I firmly believe that the earth is flat, I will rely on authorities that agree with me on this point, because I find it outrageous that some putative expert disagrees with me on this claim, and I will probably end up with many more false beliefs if I decided to rely on the authority for other questions in her domain (that the earth is a circular magnet, that nobody has ever been on the moon, and so on). The same holds for vaccinations: if a layperson allows for her object level beliefs to affect the identification of vaccines-experts, then the expected accuracy of her beliefs in the domain will depend on how accurate her core beliefs about vaccinations were at the beginning. However, if like Grundmann we do not want to allow for these judgments to guide our choice of experts, we need to be able to identify strong independent reasons to believe someone to be an expert. In the next section I will work in this direction, and build on Grundmann's proposal by arguing that, in order to provide a full account of epistemic authority, we should start from the notion of 'official experts', as the experts recognised by our institutions.

2. Official experts as epistemic authorities

In our everyday lives, we don't usually have to go around looking for reasons to believe someone to be an expert. When we are ill, we go see a *doctor*; when we learn someone is a *scientist*, that is reasons to believe they are experts; if I want to learn about quantum mechanics, I ask my *professor*. These are all social indicators of expertise: we usually don't reason about why we should trust any of these people, we just believe they are experts because of the public recognition they enjoy. Similarly, when parents need to make a decision about vaccines, they have resources available to identify experts in the field. I want to start my discussion of epistemic authorities by talking about 'official

experts', as those experts that are officially recognised to be so. Official experts are those people that are recognised to be experts by the institutions: they are doctors, professors, scientists, researchers, engineers, lawyers, and so on. They have university degrees, and other forms of public recognition accepted by the institutions of a society, such as the enrolment in professional registers, the affiliation to universities and research centres. What is the epistemic import of these social indicators, and what is their relevance for how we should respond to epistemic authorities?

Upon learning that a person is an official expert in a field, we implicitly acquire information about their epistemic position in that field⁸. First, we have reasons to believe that the official expert has acquired through her studies, training and experience, certain *competences* that are superior to those who are not official experts in that field— call these 'laypeople'. In particular, we have reasons to believe that:

- i. She has acquired a corpus of knowledge in the field during her studies, that is superior to the knowledge of those who do not have a degree (or other official certification) in the field, and includes the basic information that most people generally have from going to school or simple experience.
- ii. She has spent a longer time than laypeople thinking about the relevant questions in a domain, and she is able to employ her knowledge to answer new questions in the domain.
- iii. She has been trained to employ reliable methods to answer questions, she is able to employ the instruments that are used in her fields, she knows the protocols to follow, and she is able to reason in a way that is more reliable than laypeople.

Moreover, we learn that this person is part of a *network* of experts (colleagues, superiors, students), and this is reason to believe that:

- iv. She is in contact with other experts in the domain, and their opinions about questions in the domain (both directly, and indirectly through other experts with

⁸ For the sake of the argument, I only focus on scientific fields of expertise, and their official experts. I want to avoid the discussion about whether experts exist in fields such as ethics, politics or aesthetics. The official experts I have in mind are those for whom these conditions clearly apply. Whether they could be true for philosophers or art critics is beyond the scope of this work. However, if my analysis is correct, it could serve as a basis to discuss whether experts and authorities in these controversial fields exist.

whom she is directly connected who are also indirectly or directly connected to others). (See Kitcher 1993: 317, 318)

- v. She is part of a community of experts who assign a reputation to each other: her claims are exposed to the opinions of the other experts, and are assessed by the rest of the experts' community, either formally or informally.
- vi. Due to her contact with other experts in the field, she is in a position to acquire new evidence and to hear about new discoveries before laypeople.

Finally, the official recognition (saying that someone is a doctor, or a scientist, or a professor) indicates that a person is *active* in that field of expertise, and this indicates that:

- vii. She has access to instruments that are not available to laypeople, and this makes her superior at identifying and gathering evidence.
- viii. She is more aware of the state of the art in a field, and she is able to identify the relevant new questions in a field of inquiry better than laypeople.
- ix. Her knowledge in the field is more up-to-date than laypeople's.

Some characterisations regarding these observations are in order. First, the belief that an official expert is superior to laypeople with respect to the areas identified above is *defeasible*. If I learn that somebody is an official expert, say, a doctor, this gives me *prima facie* reasons to believe that this person is superior to me (since I am not a doctor) in all of these different things. Further information may defeat my justification for believe that this is the case: it may come out in the news that this person had a fake diploma, or that she is not up to date with the new techniques for a certain disease I am suffering from, or that she has a very bad reputation in her field. However, in the absence of counter-evidence, we have good reasons to believe that these conditions hold. Same holds for the laypeople: learning that someone is an official expert in physics gives me reasons to believe she is superior in these areas, unless I have special reasons to believe that she is not (for instance, I am a trained physicist, and I was a researcher until the day before yesterday, but I quit the job to travel around the world so that I am no longer officially considered an expert). Secondly, this superiority comes in degrees: a physics high school teacher is superior to me with regard to quantum mechanics to a lesser degree than a physics university professor specialising in quantum mechanics. So, as a physics high school teacher, you may acquire the same reasons i-ix upon learning that someone is a physics university professor specialising in quantum mechanics.

Now that we have spelled out some features that we have reasons to attach to official experts, it is time to turn to some definitions. The first relevant notion is that of *expert*.

An agent is an expert (E) for a subject (S) in a domain D if E:

- i. Given a certain body of evidence, E is more likely than S to form accurate beliefs than S on the basis of that same evidence.
- ii. E possesses more evidence in D than S.

An expert is someone who is more informed and more competent (see Elga 2007: 479). Identifying expertise with superiority makes expertise a relational concept: E can be an expert for S2, but not for S1. However, this seems to match the way the term is used: if a physics student is around a group of professors, she would not deem herself an expert. But she would if she was out for dinner with her friends, who all study psychology. A worry may be that in this way a person could be an expert for someone even though she did not *know* much herself (see Grundmann manuscript: 8; Goldman: 91)⁹. However, I don't believe this is problematic to our notion of expertise: doctors in the Middle Ages didn't have many true beliefs about the functioning of the human body. However, they were experts to laypeople, because they knew more than them. Similarly, I am an expert about physics when I hang out with my 7-years-old nephew and his friends, and if a question about physics comes up, they all turn to me to see what I have to say about that. Expertise is linked to a domain: if E is an expert in D, for every question q in the domain, E is in a better epistemic position than me to form a correct opinion about q.

An *epistemic authority* is simply an agent whose credences we adopt preemptively on the basis of *justifiedly believing that she is in a better epistemic position to answer a certain question*, ie. (i) we have reasons to believe that she is more informed and more competent on a claim, (ii) she has considered all of our relevant evidence, and (iii) she is honestly and

⁹ If knowledge is not a condition of expertise, then the existence of disagreement in a field does not threaten the existence of experts in that field (cf. Lackey 2018: 228). As I said, I am not going to focus on controversial domains such as morality and aesthetics. However, this definition of expertise could be used to argue for the existence of experts in controversial fields.

competently¹⁰ testifying about it¹¹. While expertise is linked to a domain, authoritativeness is not: I can take a person to be authoritative on specific claims, just on the basis that the conditions are satisfied in the specific case. I can take my neighbour to be an authority on the question about what time the mailman came this morning, without considering him to be an expert on the schedule of our mailman; because I wasn't home this morning so I know he has superior evidence, and the same competence in assessing it. I take expertise to be defined in an externalist sense: a person E is an expert for me if she is objectively more likely to form the correct beliefs and she has a superior body of evidence. On the other hand, authority is an internalist notion: A is an authority for me on a certain claim q, if *I have good reasons to believe that*: (i) A is more informed and more competent on q, (ii) A has considered all of my evidence relevant to determining whether q, (iii) A is reliably testifying on q. The reason why expertise is usually linked with the notion of epistemic authority, is that believing someone to be an expert in a domain means taking them to be superior in competence and information in all the questions within the domain: and this under the right circumstances gives us *prima facie* reasons to believe that they also have considered all of our relevant evidence. Therefore, if they are also reliable testifiers in the domain, experts should be treated as an epistemic authority for all the questions within their domain of expertise. However, an expert can cease to be authoritative for a claim if we have reasons to doubt either the superior competence or information or the honesty *on a specific question* (or set of questions), without ceasing to be an expert in the domain. Imagine you go see your doctor with a certain pain in the back. She examines you and she says you are fine and you don't have anything. Your doctor is an expert, so she is generally in a better position to decide about health issues, and you know from your past experience that she is a reliable testifier. Thus, you should treat her testimony as authoritative and believe you are fine. However, you may have reasons to believe otherwise if a new symptom appears of

¹⁰ By 'taking someone to be a reliable testifier' I mean that we have reasons to believe their utterances to correspond to their beliefs. Reasons to doubt the reliability of a testifier can be doubts about their honesty, or their linguistic competence.

¹¹ I agree with Grundmann that the reason for preempting our own reasons for adopting the authority's belief is that disagreeing provides us with an undercutting defeater for our belief that makes our own evidence rationally unusable. This holds insofar as we have reasons to believe that the authority has considered all of our evidence- and more, and that she is better than us at assigning right credences on the basis of that evidence. I have added the condition of being a reliable testifier, which seems crucial to attribution of authority independently of expertise.

which your doctor is unaware (evidence); or if she told you she hadn't slept at all in the past three days (competence); or if you learn that she is a compulsive liar or that she is not proficient in the language (reliable testimony). She doesn't cease to be an expert under any of these circumstances (she still has an overall larger body of evidence on health issues from her studies, past experience and so on, and superior competence at assessing it), but she is not an authority *on your case* anymore.

So far I have relied on an intuitive understanding of 'domain', but a more precise definition is needed for understanding the exact scope of experts' epistemic authority on laypeople, and laypeople's independent resources for assessing experts in specific cases, such as the case of vaccination scepticism. Following Goldman (2001: 94, 106), Grundmann defines a domain as the set of propositions that are esoteric for the layperson and exoteric for the expert. While Goldman defined an esoteric claim as a proposition whose truth-value is inaccessible to the layperson, Grundmann wants to allow the possibility for laypeople to have access to the truth-value of the proposition in a domain, and defines it as a proposition "that laypeople are *less reliable* in assessing than the authority, due to their less reliable arsenal of methods" (manuscript: 12). Thus, a domain is the set of propositions that are best assessed via specific methods. However, it is not clear why this should be so: even if they just used methods like normal inference or perception, sometimes the experts are just better and more reliable at doing that. My doctor can know that I have a cold just by looking at me, and yet I would not be as reliable if I did the same. The experts' superiority in evidence and competence explains their superior reliability. Thus, the domain should simply be identified with the set of facts on which the experts have superior competence and evidence. Just like expertise, the domain is relative to the agent for whom the expert is an expert. A doctor is an expert to me regarding medicine: no matter which methods she employs, she is generally better than me at assigning the right credence to health-related propositions. However, a paediatrician is not an expert in medicine to an oncologist (which she may be to me): she is an expert in children diseases to the oncologist, a peer in general medicine, and the oncologist is an expert to her when it comes to tumours.

The motivation for Grundmann and Goldman to restrict a domain to claims that are not accessible, or not-as-reliably accessible to laypeople comes from reflecting on cases like the following:

If Hawking predicted the visibility to the naked eye of a certain comet at time t , his belief would seem to have the weight of authority for a layperson. So it would appear that propositions like *the comet will be visible to the naked eye at t* are part of his domain of authority. But when a layperson carefully looks at the sky at t and fails to see a comet, she appears to be perfectly rational in not relying solely on Hawking's belief. The explanation for this is that the proposition at issue was part of the domain of authority before t but stopped being part of that domain at t . The truth-value of the proposition before t was assessable only via the specialized methods that allow predictions within the realm of physics, and the authority qua expert has not only mastered more of these methods, but is also more capable in applying them and additionally has superior evidence with respect to the prediction (Grundmann manuscript: 11)

However, in order to explain this case, it is not necessary to claim that the proposition *the comet will be visible to the naked eye at t* ceases at time t to be part of the domain of expertise (if it was part of the domain to begin with). The reason why this claim is not authoritative anymore, is that at time t , looking at the sky and failing to see a comet, we have reasons to believe that we are in possess of a piece of evidence (the comet was not visible) that the expert does not have. Moreover, we would have reasons to doubt that the experts' competence is superior, if the *only evidence* in favour of the truth of the proposition was the actual seeing of the comet: we would think that we are just as good as the expert to conclude: 'the comet is visible to the naked eye', upon seeing one (even though it is most likely that Hawking would have much more evidence for that claim). In any of these cases, Hawking is still an expert in the domain, he just ceases to be authoritative on that exact claim. But if Hawking was there and was looking at the sky with us, then we should probably defer to his authority even if we failed to see the comet (in the absence of other reasons to distrust him, other than our failure to see the comet). Therefore, with authority and expertise so defined, there is no reason to introduce the distinction between esoteric and exoteric claims: a domain of expertise is just the set of facts for which the expert has superior competence and superior evidence. The propositions within a domain of expertise need not be inaccessible to laypeople (à la Goldman), nor do they need to be best accessible to the expert because of their methods (à la Grundmann). This is relevant to the controversy of vaccinations and similar ones: the domain of expertise includes propositions which may be accessible to laypeople, and for which the expert may not have any superior method, just because the experts' superior competence and evidence makes them better than laypeople at assessing them.

For instance, laypeople may rely on perception to determine whether vaccines have caused a certain side effect in their children¹². Even if experts relied on exactly the same method— i.e., perception— the proposition ‘the vaccine has caused this side effect’ would still be within *their domain of expertise*, because they would still have superior competence and additional evidence (e.g., past experience, their studies, other relevant background information, the distinction between correlation and causation...).

Now, back to the official experts. I believe that the conditions sketched at the beginning of this section imply that that upon learning that someone is an official expert in a field we acquire *prima facie* reasons to believe that they are experts in that field, that they have examined all of our evidence, and they are reliable testifiers, and therefore that in the absence of defeaters for any of the three conditions (expertise, same evidence, honesty/linguistic competence) we have reasons to treat the official experts as authoritative on all the questions in their domain of expertise. The official experts’ studies, training and experience are reasons to believe that official experts possess *all of the relevant evidence of laypeople* and more. Moreover, the instruments and contact with other experts provide *prima facie* reasons to believe that even new evidence in a field reaches the experts before laypeople. Their ability to identify and gather evidence, together with their methods and instruments, and the community of which they are part are reasons to believe that the official experts possess a *substantive body of evidence*, superior to that of laypeople— even the most informed ones. Their training and experience are reason to believe in their *competence*: they studied for many years, passed exams, trained their skills, learnt the most reliable methods. Finally, the network structure of official experts gives reasons to believe that they are reliable testifiers: the reputation that experts assign to each other both formally and informally is a reason for laypeople to believe that the reliability of experts is monitored: charlatans are expelled from public registers, research is subject to peer review, conflict of interests are monitored. Without reasons to doubt the sincerity of an expert on a specific claim, the layperson has *prima facie* reasons to believe they are reliable testifiers. Therefore, in the absence of reasons to think otherwise, the fact that an agent is an official expert is *prima facie* reason for a layperson (i.e., someone who is not an official expert in the same field) to treat them as epistemic authorities on all the questions within their domain of expertise. Of course, this holds

¹² There is a paper, by the EFFV, which reports all the cases of side effects of vaccines in which the causal relation had been denied by the experts. Available here: https://www.efvv.eu/wp-content/uploads/2018/03/EFVV_report.pdf

for the case of vaccinations. In the absence of reasons to think otherwise, we have reason to consider *every official expert* in a domain that includes vaccinations (for instance, medicine) to be an epistemic authority for us *for every question* on the topic of vaccination. Hence we should preemptively adopt their beliefs on the matter, simply on the grounds that they are epistemic authorities.

Of course, official experts are not the only people we have reasons to treat as authorities. However, especially in the case of highly specialised domains of expertise, where laypeople may not have much understanding, it may be very hard for laypeople to have good reasons to believe that the conditions for authority are fulfilled. Without the possibility of relying on social indicators of expertise, it may be very hard for me to determine whether an agent is an expert in, say, astrophysics, knowing that my knowledge in astrophysics is very limited. How could I acquire reason to believe that someone is an expert and they have considered all of my evidence and they are reliable testifiers without relying on social indicators of expertise? It is easier to have reasons to treat someone as an authority on specific claims, since it's easier to check that the conditions have been satisfied. It is surely more difficult to assign authority in a whole domain, especially for those domains in which our own epistemic position is not very favourable and where we don't have the possibility to verify the putative experts' claims. One such domain is vaccinations: how could I acquire evidence about another layperson being an expert, having considered all of my relevant evidence on the topic of vaccinations, and being a reliable testifier, without relying on social indicators of expertise (such as a degree, or an institutional recognition)? This would require a great amount of additional information regarding the putative expert, which, as laypeople, we are hardly in the position of acquiring. However, it is a possibility that I am not excluding, and I shall come back to this later in the paper.

3. Dependence on the community of official experts for coverage

In the previous section I have argued that we have *prima facie* reasons to treat the official experts as authoritative for every claim they make within their domain of expertise. For instance, parents have reasons to treat their paediatricians as epistemic authorities on the topic of vaccinations, and to follow their advice in a preemptive manner. This is because the features that are associated with official expertise give us good reasons to believe that official experts (i) are experts, (ii) have considered all of

laypeople's relevant evidence, (iii) are reliable testifiers. These conditions may be met by people other than the official experts, but in those cases we would need to gather positive reasons to believe so, and given the epistemic asymmetry that characterises the relation between experts and non-experts, it may not be easy to obtain them. Providing good *prima facie* reasons to be treated as authorities in a domain is not the only special feature that official experts enjoy as opposed to other putative experts. Another kind of epistemic reliance that is *prima facie* justified in the case of official experts is the reliance on their lack of expressed opinion. Intuitively put, it seems that for official experts the following *truth-to-testimony* conditional (Goldberg 2010a; 2010b; 2011) holds: if *p* were true (where *p* is a claim within an experts' domain) an official expert would have reported about it by now. For instance, it seems rational for me to think that it can't be the case that baking soda cures cancer, because if that were true I would have heard about it from an official source¹³. In this section I shall introduce and explain the truth-to-testimony conditional as defended by Goldberg (2010a; 2010b; 2011) and the *absence-based* beliefs that are justified by it, and illustrate why it is a valid inference with respect to communities of official experts, but not for other communities of laypeople (including unofficial experts).

The inference from absence of evidence is generally considered a fallacy: absence of evidence is typically *not* to be taken as evidence of absence. However, some version of the argument from ignorance has recently been defended in argumentation theory (Tuzet 2015), cognitive science (Hsu et al. 2017), and epistemology (Goldberg 2010a, 2010b; 2011; Pedersen & Kallestrup 2013). Even though the truth or falsity of a claim cannot be *deductively* inferred by the lack of evidence in favour or against it, it is under certain circumstances reasonable to inductively or abductively infer the truth of a claim from the absence of evidence against it (Tuzet 2015: 39). Lacking evidence in favour of a claim does not logically imply the falsity of this claim. This much is trivial. However, in some cases, the absence of evidence in support of a claim can make very plausible the falsity of that claim. Let's consider an example: the fact that walking around a forest for the first time I have not encountered any dears is clearly not evidence that there aren't any dears in the forest. However, the fact that after years of monitoring the woods the forest rangers have not found any dears *is* evidence that there aren't any dears in the forest. The

¹³ An official expert did in fact report this, but he stopped being an official expert shortly afterwards: http://www.ansa.it/english/news/general_news/2018/01/15/doc-gets-5-yrs-for-treating-cancer_5a91a283-1572-4a99-ac63-5653fcdf5b3a.html

crucial factor distinguishing these two cases, and making the latter inference from absence of evidence justified, seems to be related to the thoroughness and completeness of the search (see Tuzet 2015: 43).

In epistemology, the topic of inference from absence of evidence has been investigated by Sanford Goldberg as a form of epistemic reliance on others: the reliance on our environments for *coverage* (Goldberg 2010a, 2010b; 2011)¹⁴. The focus of Goldberg's analysis is what he calls coverage-supported beliefs, i.e., beliefs that are sustained by appealing to the truth-to-testimony conditional 'if p were true I would have heard about it by now, therefore $\neg p$ ' (2010a: 158). Coverage supported beliefs are beliefs in the truth (or falsity) of a claim inferred from the fact of not having encountered any pieces of testimony against (or in favour) of that claim under a certain set of circumstances. For instance, since I check the Italian news every morning, I justifiably believe that the new Prime Minister has not been chosen yet, for *if that was the case, I would have heard about it by now*. Being connected to my epistemic environment in an appropriate way as far as Italian politics is concerned, I justifiably take the absence of evidence (the fact that I haven't encountered any reports on the matter) to constitute evidence of absence. Goldberg identifies three conditions that need to be met in order for absence-based beliefs to be justified.¹⁵ The first condition is the *source-existence* condition, according to which in order for a belief to be coverage supported there needs to be a source in the hearer's community disposed to make the relevant facts available to the hearer. The second condition he identifies is the *reliable-coverage* condition, which requires that the source will reliably investigate the questions in her domains, and will reliably report about them 'in a timely fashion' (Goldberg 2010a: 160). Finally, the *sufficient-interval* condition requires that for the hearer's expectations to be justified, there needs to be an interval of time sufficient for the source to made available any discoveries. That is, at the time the hearer relies on the source for coverage, it must be the case that if there were any discoveries to be made, the source would have discovered them and reported about them. In addition to these three conditions on the source, according to Goldberg, it must be the case that the hearer is *receptive* to eventual reports of the source (2010a: 164). If,

¹⁴ Goldberg has made a similar case for memory-based beliefs, but I here I shall only focus on his case for testimony-based beliefs. For the memory case see Goldberg 2010b.

¹⁵ Being a reliabilist, this for Goldberg means reliably formed.

despite the receptiveness, no reports are encountered by the hearer at a certain time, and the three other conditions hold, then the hearer's absence-based belief is justified.

A more precise formulation of these conditions has been provided by Pedersen and Kallestrup (2013). Their paper is a clear summary and formalisation of Goldberg's intuition, particularly useful to our purpose of applying this discussion to the case of experts and authorities, and includes an evidentialist justification of absence-based inference that was lacking in Goldberg's original work¹⁶. Pedersen and Kallestrup unify the three conditions that apply to the source in one threefold condition, the *Epistemic Coverage Condition*:

(ECC) At a given time t , source S^* provides epistemic coverage for subject S within domain D just in case:

- (i) S^* reliably tracks p -relevant evidence and reliably determines whether p on that basis (for p that S takes an interest in and pertains to D),
- (ii) S^* is reliable in making p -relevant evidence available to S , and
- (iii) if relevant p -evidence were to be available, S^* would track it and make it available to S by the time at which S relied on S^*

(Pedersen & Kallestrup 2013: 2578)

The epistemic coverage condition requires that the source reliably investigates in a domain, and communicates eventual discoveries in a timely fashion. In addition to ECC, they reformulated the condition on the hearer, which they name *Receptivity Condition*:

(RC) If p -relevant evidence were to be made available by source S^* on which the subject S is relying, then S would register it. (Here S^* is a source that provides epistemic coverage for S relative to p 's domain.)

(Pedersen & Kallestrup 2013: 2581)

The justifiedness of an absence-based inference is the result of both the reliable coverage provided by the source, and the appropriate receptiveness of the hearer. The

¹⁶ Pedersen and Kallestrup rely on an evidential counterpart of the alethic Equivalence Schema. Without going into the details of it, the result of their analysis is that, for a restricted domain, in which ECC, RC and AC are satisfied, the alethic Equivalence Schema and the evidentialist Equivalence Schema come relatively close, and a high proportion of the evidential instances of the schema come out true. Thus, given the epistemic coverage of a source and the receptivity of a hearer, 'absence of evidence sufficiently good to justify belief amounts to evidence sufficiently good to justify belief in absence' (Pedersen & Kallestrup 2013: 2592).

third condition regards S's lack of received evidence in support of p. However, this evidential absence needs characterisation— Pedersen and Kallestrup call this the *Absence Condition*:

(AC) Enough time has passed for source S* to track evidence supporting p and make it available to S, but no such evidence has been made available to S by S*. (Here S* is a source that provides epistemic coverage for S relative to p's domain.)

(Pedersen & Kallestrup 2013: 2581)

If these three conditions are satisfied in a domain, then the inference from absence of evidence (so characterised) to evidence of absence, is valid and the corresponding absence-based belief justified.

What can the discussion on absence-based inference contribute to our discussion of epistemic dependence on experts? I believe that the community of official experts, as a whole, provides something similar to epistemic coverage for laypeople in every domain of official expertise. Having defined an expert as a superior both with respect to competence and evaluation, in the last section I defined a domain of expertise as the set of facts for which an expert is epistemically superior. If we consider all the domains in which there are official experts, we have reasons to believe that in all these domains official experts will be epistemically superior to laypeople both regarding evidence and competence. Moreover, I believe we have reasons to consider official experts as generally reliable in their fields: the technologies developed through science, the progress we have made in all scientific domains, the predictions made by scientific experts which can be verified by laypeople leave very little room for radical scepticism about science. It seems reasonable, for now, to grant that official experts, taken together as a community of experts, are generally reliable in their investigations, and not just superior to laypeople.¹⁷ From the characteristics of official experts I have sketched in the previous section, there are two more observations that seems to be justified for official experts. First, official experts are in a position to identify relevant questions in their domain better than laypeople. The larger body of evidence possessed, together with their methods, and the connectedness in their domain of expertise, make them more apt at identifying the questions worth investigating in their domain. Secondly, it follows from the way in which scientific research is structured that researchers have incentives to look for novel interesting questions, and to publish any interesting positive result. These observations

¹⁷ I shall come back to this point in the next section.

ground the attribution of ECC to official experts in each domain, taken together as a community:

- (i) The community of official experts in domain D reliably tracks p-relevant evidence and reliably determines whether p on that basis. (Where p is any question within the domain that is deemed worthy of being inquired by official experts in that domain)
- (ii) The community of experts reliably makes any results of their investigations available if it meets the standards of good scientific practice (by means of peer-reviewed publication).
- (iii) Were discoveries to be made regarding p, they would be made available by the official experts by the time laypeople relied on them.

These three conditions taken together are not yet sufficient to justify absence-based belief, until more considerations regarding the timeliness of the inquiry. However, for now this is sufficient to justify the weaker claim that in the absence of official experts' reports for a claim within any domains of official expertise, it is *irrational for a layperson to assign any positive credence to that claim*.

The absence of official experts' reports on the truth of propositions within their domain of expertise can be the result of three different scenarios: it may be that there is an ongoing investigation trying to determine whether p, but evidence has not yet been found; or, it could be that an investigation has been done, but no significant result has been found that establishes p, and the official experts believe $\neg p$; or it could be that no investigation is being done on whether p, because the experts believe that the question whether p is not worthy of investigation, either because they believe that p is false or because they believe that it is at this moment impossible to establish whether p. If some version of the receptivity condition¹⁸ is fulfilled, the lack of official reports encountered by the layperson is evidence that one of these three situation is the reason. Therefore, were the layperson to adopt a positive credence in p, she would be in the situation of

¹⁸ The RC requires that the subject would track reports regarding p if the source made them available. In the case of experts and laypeople this condition would be satisfied if the layperson thoroughly looked for reports herself (for instance by searching in the peer-reviewed journals), or if the reports were diffused beyond the groups of experts to wider audiences via media that the subject reliably monitored for new information (for instance, scientific magazines). Axel Gelfert has referred to this property of information of being diffused into one's social environment as epistemic penetration (Gelfert 2013: 781).

disagreeing with someone who is an expert, has considered all of her evidence and is a reliable testifier. This in turn, as we have seen in the previous sections, would defeat whatever reason she may have had to assign a positive credence to *p*. The absence of official experts' reports regarding *p* is an *undercutting defeater* for a layperson's belief that *p*, where *p* is any claim that belongs to any domain of official expertise (which we have restricted to scientific domains of expertise). This is a first way in which laypeople should rely on experts' silence, in addition to relying on their testimony. Moreover, if a layperson has reasons to believe that an investigation occurred regarding *p*, and enough time passed for the experts to make any findings available, this is reason for the layperson to believe *p* to be false. This is the sort of reliance described by Goldberg. In any case, the silence of experts bears epistemic significance for laypeople, and, in the absence of reasons to doubt that any of the conditions above are fulfilled, it is never permissible for a layperson to believe a domain-relative claim for which there is no official report. For instance, if no official expert has reported about a certain side effect of vaccines, or about the inefficacy of vaccines, it is irrational to believe any of these claims. Moreover, if enough time passed and we have reasons to believe an investigation has occurred (as, for instance, in the case of MMR vaccines and autism¹⁹), this is reason to believe these claims to be *false*.

Let's recap what our argument has been so far. I have proposed an account of epistemic dependence on experts which is grounded in the way in which we typically encounter and identify expertise in our lives, through official experts. The focus on official expertise has proved fruitful in many respects. First, it has allowed us to make the discussion about epistemic authority more concrete: we focussed on the ways in which we actually encounter experts in our lives, and on the characteristics we have reasons to attribute to these experts simply by thinking about how expertise is acquired and structured in our societies. Secondly, an analysis of the notion of official expertise brings together the two debates about how to respond to experts and how to identify them. The discussion of epistemic authority in the current debate is characterised by a conditional form: it is a debate regarding how it is rational to respond to epistemic authority on the premise that one recognises someone else as being in a better epistemic position to answer a certain question. However, this ignores the fact that it may be very hard for laypeople to reliably identify experts, or even know what are good reasons to believe

¹⁹ I will discuss this example in more detail in section 6.

someone to be in a better epistemic position. And this in turn has motivated some critics of the preemption view to object that the preemption view may lead to paradigmatically irrational beliefs which appear completely justified to the subject believing them on authority (Lackey 2018: 234). By shifting the focus to official experts we managed to analyse together the rationality of believing on authority with the rationality of believing someone to be an authority, especially in those domains where laypeople may not be in a position to acquire independent reasons to believe someone to be an authority. Finally, by focusing on official experts, we have shed light on another dimension of epistemic dependence on experts, namely, the reliance on the community of experts for coverage. The dependence on official experts' silence could not be derived simply from the definitions of experts and authorities, and is based on an analysis of the actual conditions in which official experts operate. The fact that the community of official experts provides epistemic coverage for laypeople in every domain of expertise does not only depend on their superiority, but also crucially on the actual ways in which research is carried out in our societies, on the incentives that researchers have to find interesting questions in their domains, and to make the results of their investigations public.

At this point, we can draw our first conclusions for the case of vaccination scepticism. In the absence of reasons to think otherwise, we have reasons to trust what the official experts say about vaccines, and to mistrust anything that has not been claimed by official experts. Moreover, we ought to do this preemptively: official experts' opinions (or lack thereof) are not *additional* evidence that we ought to consider together with other reasons we may have to trust or mistrust vaccinations. Official experts' claims and silence provide a reasons to believe or disbelieve a claim that screen off any other relevant reasons we may have. In other words, as laypeople, we *should not* engage with the actual arguments and evidence for and against vaccines, and weight our own reasons derived from this assessment against the additional evidence that experts believe this or that. Rather, as long as we have reasons to believe that someone is an epistemic authority for a certain set of claims, such as vaccines, we should simply trust what they say, on the basis that they say it. However, until now, I have relied on an idealised picture of official expertise, and left aside specific possible complications that may lead laypeople to be sceptical of their official experts. In the next three sections, I shall investigate under what circumstances it is reasonable for laypeople *not* to defer to the authority of official experts' claims and silence, and what sources of evidence are available to laypeople trying to decide in the controversial cases, without the necessity of acquiring expertise in the field. I shall

consider the issue of disagreement between experts, and other reasons to doubt that the conditions of epistemic authority or epistemic coverage are fulfilled. I shall then propose heuristics that may help laypeople decide which experts they can trust in these cases.

4. How can laypeople assess conflicting authorities' reports?

I argued that official expertise is a reliable way to identify epistemic authorities in a domain. Given the way in which official expertise is attributed and monitored in our societies, we have *prima facie* reasons to believe that official experts are epistemic authorities— i.e., (i) are experts, (ii) have considered all of laypeople's evidence, (iii) are reliable testifiers— for every claim they make within their domain of expertise. However, a layperson may acquire further evidence that casts doubts upon the fulfilment of one or more of these conditions in a given case. For instance, it can be the case that the layperson acquires evidence to doubt the legitimacy of the attribution of official expertise to a certain individual: the layperson may acquire evidence that despite the social recognition that some official expert has, the title does not in this case indicate that the features of official expertise are satisfied (e.g., the official experts' degree was a fake). Even though the official expert enjoys the reputation of being an expert, in this case the official expert just ceases to be such to the layperson's eyes. However, it may also be the case that the layperson doubts the authoritativeness of the official expert on a claim or set of claims within their domain, while still believing the veracity of the official experts' expertise, if the layperson has reasons to doubt the competence and information of the official expert on a specific topic, or their reliability as testifiers, or if the layperson has reason to believe that the official expert has not considered a certain piece of relevant evidence. In such cases, the layperson may still consider the official expert an expert, given their overall superior competence and information, but they would cease to regard them as authorities on a specific claim or set of claims. These localised cases of scepticism of the official experts' authority can be easily solved by laypeople by consulting *another official expert* who is immune to the same scepticism. For instance, imagine a layperson, Anna, who needs to decide whether to get the HPV vaccine. She goes to her old family doctor, who tells her that the vaccine is effective, and the side effects mild and rare. However, she worries because a friend told her about a new study showing potential deadly side effects, and she has reasons to believe her doctor has not heard about it, because he told her that he did not attend any conferences or refreshment

courses on the topic in recent years. What should Anna believe? Since she is not an expert, she knows that any official expert would be more reliable than her at forming a judgement on the matter. However, she has reasons to believe her doctor lacks a piece of evidence that she has. The rational thing to do for her is to go see another doctor, maybe a specialist, who is more updated in the field and has surely heard about any new studies.

The main reason for vaccination scepticism, though, does not seem to be related to doubts about individual experts, nor to mistrust in experts in general. It only takes a brief internet search to find yourself swamped by information about how vaccines cause autoimmune diseases, or how they cause aluminium to settle in brain tissue, and other effects. Some of these reports are not well supported, are not produced by official experts and have only vague references to studies of experts. However some *are*. The anti-vaccination movement has supporters among the experts, and some of the reports about the health risks of vaccination come from doctors and scientists. To be clear, vaccinations *do* have health risks, and this is acknowledged by all the parts involved. However, the experts seem to *disagree* on what exactly these risks are, on their magnitude and occurrence, and on whether the risks of vaccinations are clearly outweighed by their benefits. The problem faced by laypeople, and the reason that may lead parents to be hesitant towards vaccinations, is that some official experts *disagree* with the most part of the other official experts, who support vaccinations. How are laypeople supposed to decide what to believe when official experts disagree?²⁰

Official expertise, as we have seen in the previous sections, is a good guide to identifying epistemic authorities. However, identifying epistemic authorities is not the only problem faced by laypeople in their dependence on experts: experts often disagree with one another, and in these cases it is necessary for laypeople to *assess* experts, in order to decide whom to trust. To some extent, official expertise may also be a guide to assessing expertise: official expertise comes in degrees, and it may be possible for laypeople to identify a hierarchical structure within official experts by looking at their titles and roles. However, the potential of official expertise for assessing experts is only limited: first, it may be that an official expert disagrees with a non-official expert for whom the layperson has good reason to believe that all the conditions of authority hold,

²⁰ The problem occurs when epistemic authorities disagree in general, not necessarily official experts. However, as I have claimed before, it is very hard to see how a layperson may acquire reasons to believe another layperson satisfies the conditions of epistemic authority in fields such as vaccinations, without relying on social indicators.

but no social indicators of expertise are present; secondly, two experts that are officially recognised as peers may disagree; third, even when two official experts who are not officially peers disagree, it is not necessarily so that the layperson should treat as an authority the one with best credentials, given that the conditions of epistemic authority include factors that are not necessarily related to superior official credential (e.g. honesty). When two experts who satisfy the conditions of epistemic authority when individually considered, disagree on a claim or set of claims, this constitutes evidence that one of them is mistaken on that claim. How can a layperson decide which of them is to be treated as an authority with regard to the claim at stake?

The first thing to notice is that it is perfectly consistent with what we have said so far that a layperson may not adopt the belief of someone who, when considered in isolation from other experts, would satisfy the conditions of epistemic authority. One of the objections put forward against the preemption view of epistemic authority is that a layperson caught between disagreeing epistemic authorities— call them S and S*— faces the following dilemma:

If you preemptively adopt S's position, you cannot at the same time rationally adopt S*'s position, and vice versa. In either case you would thus hold a belief that contradicts a belief held by someone you conscientiously judge to be an epistemic authority. On the other hand, if you preemptively suspend belief then you adopt an attitude that conflicts with the attitudes of both of your authorities. (Jäger 2016: 172)

Jäger's critique may be spot on for Zagzebski's account of epistemic authority. However, it loses its force given the definition of epistemic authority we are operating with. An epistemic authority is someone whose beliefs we adopt preemptively on the basis of judging them to be in a better epistemic position than ourselves to answer a certain question. In particular, an epistemic authority is someone whom we believe to be more informed and more competent on a certain matter; has considered our relevant evidence on the matter and is reliably testifying about it. However if we encounter two people for whom these conditions hold, and yet they disagree with each other, this counts as evidence that one of the two is mistaken in their assertion. In the absence of reasons to

prefer one point of view over the other, the rational thing to do is to suspend judgement.²¹

The crucial question now is: what would be a good reason to prefer one point of view over the other? There are two possible candidates: on the one hand a layperson may think that she should go with the point of view that she deems more plausible, given the evidence that she has; on the other hand, the layperson may try to assess the rival expert's authority on the claim indirectly, through an assessment of their expertise. I believe that the latter is the only rational strategy among the two. When a layperson is faced with the disagreement of two experts who individually taken satisfy the conditions of epistemic authority, and she has no reasons to believe that they are not peers, then the disagreement is in her eyes a disagreement among two individuals who (i) have superior information and competence than her, (ii) have considered all of her evidence, (iii) are reliable testifiers, and yet hold conflicting beliefs regarding the matter at stake. If she were to employ her own evidence to decide which one holds the most plausible view, she would be behaving irrationally since she believes that *both* have already considered her relevant evidence on the matter. If she decided to treat the first as an authority, she would be irrationally disagreeing with the second expert, who has considered her evidence, and has come up with a different opinion. It would be different if the two experts who satisfy the conditions of authority were not peers: in that case, she would be justified in following the expert who is in the better epistemic position to answer the question at stake due to their superior expertise, just as they are justified in disagreeing with the other expert. In order to rationally decide between the two rival authorities she needs to acquire evidence *that they are not peers* (or that they don't both satisfy the conditions of epistemic authority). However, given her inferior competence in the domain, she cannot establish this by judging their claims in the domain: in order to

²¹ To the layperson's eyes, this situation is one of peer disagreement. According to many philosophers, peer disagreement calls for a suspension of judgement. Richard Feldman, for instance, writes: 'In those cases, I think, the sceptical conclusion is the reasonable one: it is not the case that both points of view are reasonable, and it is not the case that one's own point of view is somehow privileged. Rather, suspension of judgement is called for'. (2006, p. 235) This does not mean that to the experts who are disagreeing this is a situation of peer disagreement. The experts don't necessarily believe that they share the same evidence and the same competence, so they don't need to suspend judgement; nor they need to defer to each other, if they don't have reasons to believe the other to be superior with respect to both competence and evidence, and to have considered all of their evidence, and to be reliably testifying.

decide whether they are in fact peers she would have to be reliable at discerning true and false propositions in the domain, which she is not. Hence, in order to decide which of two competing authorities to trust on a specific claim or set of claims, the layperson needs to look at indirect indicators of expertise, where superior expertise (with the other conditions being equal) is related to a higher likelihood of being correct on the claims in the domain of expertise, including the controversial claim.

Alvin Goldman identifies five potential sources for non-experts to indirectly evaluate experts' conflicting testimony:

- (A) Arguments presented by the contending experts to support their own views and critique their rivals' view.
- (B) Agreement from additional putative experts on one side or other of the subject in question.
- (C) Appraisals by "meta-experts" of the experts' expertise (including appraisals reflected in formal credentials earned by the experts).
- (D) Evidence of the experts' interests and biases vis-à-vis the question at issue.
- (E) Evidence of the experts' past "track-records". (Goldman 2001: 93)

According to Goldman, this is a (non-exhaustive) list of possible ways in which laypeople may decide among conflicting authorities without directly engaging with the arguments and evidence for the claims they make. To adapt his list to my terminology²²: methods (A) and (E) are ways of assessing indirectly which of the competing authorities has superior expertise, by looking at the ways in which they present their work, and their performance in the past. Method (D) is not a way of assessing expertise, but rather possible evidence that one of the experts is not a reliable testifier, and should not be treated as an epistemic authority. I shall discuss these methods and their limitations in the remainder of this section. I leave the discussion of the methods (B) and (C) of additional support by experts and meta-experts to the next section.

Let's start with argumentative superiority. If the layperson trying to assess two rival experts on a claim lacks expertise in the field, then she will probably not be in the position of evaluating the argumentative superiority of an expert over the other directly, i.e., by judging the premises of the argument and their relation to the conclusion.

²² Goldman has a different notion of expertise, domain and authority (Goldman 2001). However, his methods can be easily adapted to our discussion.

However, according to Goldman, indirect evaluation of argumentative superiority may be possible, in the form of *dialectical superiority*: the ‘argumentative performance’ (Goldman 2001: 96) of an expert may be an indicator of greater expertise. What Goldman means by dialectical superiority is not merely the debating skills of a putative expert, but rather the way in which she presents counterarguments to the opponent, whether she seems to have considered the relevant objections, and whether she can rebut them quickly and smoothly. The layperson can acquire indirect argumentative justification to trust an expert over another by an inference to the best explanation from the argumentative performance to the level of expertise of an expert. However, there are some limitations to argumentative superiority as a source for expertise assessment. First, as Goldman himself acknowledges, this assessment is only possible when the layperson witnesses a full debate in which the rival experts engage in a discussion and detailed defence of their positions against the opponent (Goldman 2001: 94). However, laypeople are rarely in a position to witness such discussions, and most often only encounter sketches of the expert’s arguments for a theory, while the full defence and argumentation is reserved to a restricted audience. Moreover, the reliability of dialectic superiority as an indication of expertise is questionable, and so is the ability of laypeople to assign the right weight to the right features in experts’ debates: a lack of rebuttal may be due to an expert’s judging an objection to be not worth of response, or to having been answered in another occasion, and laypeople as such are usually not in a position to tell which arguments are relevant, and what the status of the dispute is between experts (see Almassi 2012: 36). It seems that in order to be able to reliably evaluate an argumentative performance as an indicator of expertise more needs to be taken into account than the pure formal dynamics of the dispute. As Almassi notices, even the simple indicator of self-contradiction becomes controversial, if it is laypeople who need to judge whether an expert has contradicted herself or not (2012: 35). Not being an expert in a field seems to involve more than just not being able to engage with the content of an expert’s claims: even when laypeople do witness a full debate between rival experts, the lack of expertise often involves not being in a position to distinguish a good argument from a bad one (not even formally), not being able to distinguish a relevant objection from one that does not deserve being answered, or not being aware of the status of a discussion between rival experts²³.

²³ On the topic of experts debating with an audience of laypeople, and the complications that arise in this kind of argumentation, see Mercier (2011: 321, 324).

A similar difficulty holds for the two sources of evidence of biases and interests (D) and evidence of the experts' past track records (E). If one acquires evidence of the interest of a speaker in reporting a certain fact that is misaligned with the truth of the report, this surely counts against the speaker's credibility. Similarly, if the reports of two speakers conflict with each other, and only one of the speakers has such interests, in the absence of other reasons to prefer one over the other, it seems reasonable to prefer the report of the uninterested speaker. However, for experts, evidence of this kind may be very hard for laypeople to obtain and to interpret correctly: often experts have different interests that compete with each other, and these interests are not necessarily misaligned with the truth. Official experts have personal interests in improving their reputation, in advancing their career, in obtaining the respect of colleagues, and so on. These interests compete and often prevail over other possible interests, and are mostly aligned with finding and reporting about the truth. In some very obvious cases laypeople may reliably judge among experts on the basis of their interests. However, it may sometimes be hard for laypeople to identify conflicting interests and decide on this basis which experts are more reliable. The same holds for biases, with the further complication that biases are even harder to detect than interests, are widespread among cognitive agents and very hard for individuals to get rid of (see Trout 2004). With the exception of some very obvious cases, where it is clear that the experts' testimony was not guided by the truth, evidence of biases and interests will not be of much use to laypeople trying to assess experts.

What about evidence of the experts' past track records? Here, it's helpful to follow Thi Nguyen's distinction between *obvious domains*, *subtle domains*, and *cognitive islands* (Nguyen forthcoming). Obvious domains are domains in which the experts' actions have an impact on the world outside of the expert's domain (either by making predictions, or by producing technological devices), such that laypeople can assess the failure or success by means of observations that do not require any expertise (Nguyen forthcoming: 5). An example of an obvious domain is aeronautical engineering: non-experts may be unable to assess individual claims made by aeronautical engineers regarding planes, but they can assess their track records by looking at the fact that planes fly. For subtle domains, on the other hand, there is no such tests available. Subtle domains are domains in which significant expertise is required to tell the good cases from the bad (Nguyen forthcoming: 7). An example of a subtle domain is particle physics. Even though non-experts do not possess the expertise to assess expertise in subtle domains, sometimes

these domains are connected to obvious domains and can therefore be indirectly assessed (Nguyen forthcoming: 8): nuclear engineers rely on particle physics, and the failures and successes of nuclear engineering may be accessible to laypeople. When a domain is both subtle and does not bear connections with obvious domains, it is a *cognitive island* (Nguyen forthcoming: 9). Lacking any connection with accessible claims, cognitive islands cannot be assessed by laypeople. Even though Nguyen's examples of cognitive islands are the moral and aesthetic domain, some empirical disciplines too can be considered cognitive islands—at least for a certain period of time²⁴. Thus, evidence of an experts' past track records is only open to laypeople's assessment in obvious domains and, with more difficulty, in subtle domain (of course, it becomes increasingly difficult as the connection between the subtle and the obvious domain becomes more distant)²⁵. While the field of medicine, in general, seems to be an obvious domain, this only supports the idea that laypeople should *prima facie* trust the official experts in the field. However, this is not of much use when experts in the field disagree, as it happens with vaccines. It is true that vaccines have (almost) eradicated several diseases, and keep many other under control, which should give confidence in the vaccines experts, but what is doubted by vaccination sceptics—vaccines experts included—is whether many side effects have been taken into account and researched sufficiently, and whether it is still necessary nowadays to perform so many vaccines. With regard to these issues, the general track record of the discipline as a whole does not seem of much help. Moreover, another difficulty with track records is that laypeople may not always be in a position to gather such track records for individual experts—rather than for the discipline as a whole (see Guerrero 2016: 15). Finally, their lack of expertise makes laypeople susceptible to being deceived by experts using their track record in a field as an indication of expertise in another field, or even using it for malicious purposes, such as serving group interests, as it may happen for instance with climate change sceptics (see Martini 2014: 13; Kitcher 2010).

²⁴ Elijah Millgram, for instance, proposes the field of high-end finance as an example of a domain inaccessible to non-experts (2015:38-39).

²⁵ Even though the track record may have limited application for expertise assessment, this observation still has another important implication. The positive track record of most scientific fields (both obvious and subtle) should be enough to ground the assumption that official experts are not only superior to laypeople, but also generally reliable. We can thus, at least for the purpose of this paper, ignore the possibility of laypeople's global scepticism about official experts. For a discussion of this kind of scepticism and the comparative superiority of science, see Goldman (1999), ch.8.

Among the methods that Goldman proposes for laypeople to assess expertise, argumentative superiority, evidence of interests and biases, and past track records have limited application, and are faced with significant difficulties. In the next section I shall discuss the two methods that employ numerical aggregation: the additional support from other experts, and the appraisal from meta-experts.

5. Experts disagreement: going by the numbers

How should laypeople decide between conflicting authorities? What should we believe about vaccines safety, after being presented with conflicting opinions on the matter? Among the methods proposed by Goldman, the most promising seem to be the methods that are based on numbers: (B) the agreement of other experts with one or the other side of the dispute, and (C) appraisals of the experts' expertise by reliable third parties, the 'meta-experts'. Meta-experts are people who are reliable at identifying experts. Hence, the support of meta-experts is an indication of superior expertise that should provide reasons for the layperson to believe that the two disagreeing experts are not peers, which in turn justifies the non-expert's decision to treat one of the two as an authority on the question at stake. Goldman takes this source of evidence to include credentials, such as academic degrees and professional accreditations. If two experts disagree on a claim, and one expert has more support by meta-experts than the other, then with the other conditions being equal, this is evidence that the first expert is a superior expert, and is therefore more likely to be correct. Differently from methods (A), (C) and (E), the agreement of additional experts on the claim at stake is not a way of assessing the competing authorities' expertise, but rather a way for the layperson to assess the likelihood of that claim being true, regardless of which of the two expert making the claim in the first place is the superior expert, without employing domain-sensitive evidence. Consensus seems in general to be evidence of a claim's truth. However, Goldman considers two cases in which 'going by the numbers' does not seem a rational strategy to assess the credibility of testifiers. The first case is the case of a guru, whose followers uncritically reflect whatever the guru says (Goldman 2001:98). In this case, the number of people who agree with all of the guru's claims does not seem to bear on their truth. Another example is the case of rumour: if someone hears a rumour from a source, and passes it on to other hearers, the rumour does not become more likely to be true the more it is repeated by new hearers (Goldman 2001: 99).

According to Goldman, in order for numbers to have evidential weight in assessing the credibility of rival testifiers, the beliefs of the people who agree with each other need to be at least partially *independent of one another*. Non-independent opinion-holders agreeing with one of the parties involved, should not affect the assessment of the testifier's claim. The concept of non-independence is expressed by Goldman in terms of conditional probability. He considers two opinion holders X and Y, and their credence in a certain hypothesis H being true X(H) and Y(H). Y's belief that H is totally non-independent of X's when Y is as likely to believe what X believes if H is true and if H is false. Formally put:

$$P(Y(H) | X(H) \& H) = P(Y(H) | X(H) \& \neg H)$$

This is what happens in the case of a blind follower of a guru, or in the case of rumour. When this is the case, Y is called a *non-discriminating reflector* of X, and his opinion has no evidential weight for a subject trying to assess X's credibility (Goldman 2001:101). In order to have an evidential boost from Y's opinion, the novice should justifiably believe that X and Y are at least partially *conditionally independent* of one another:

$$P(Y(H) | X(H) \& H) > P(Y(H) | X(H) \& \neg H)$$

In other words, for Y's agreement to count in favour of believing H, the non-expert should justifiably believe that Y would be more likely to share X's opinion that H if H was true than if H was false. The reason for Goldman's conclusion comes from a simple Bayesian proof (2001: 100-101). How should a subject S update her belief that H, given that an expert X believes H? According to Bayes' theorem, the layperson should update her belief according to this 'likelihood quotient':

$$(1) P(X(H) | H) / P(X(H) | \neg H)$$

If two experts X and Y believe H, then the likelihood quotient is:

$$(2) P(X(H) \& Y(H) | H) / P(X(H) \& Y(H) | \neg H)$$

Which according to probability calculus is equivalent to:

$$(3) P(X(H) | H) P(Y(H) | X(H) \& H) / P(X(H) | \neg H) P(Y(H) | X(H) \& \neg H)$$

Now, Goldman notices, if Y is a non-discriminating reflector of X, then anything believed by X will also be believed by Y, no matter whether H is true or false. In the non-discriminating reflector's case:

$$(4) P(Y(H) | X(H) \& H) = 1 \text{ and}$$

$$(5) P(Y(H) | X(H) \& \neg H) = 1$$

This makes the likelihood of H occurring given X's belief that H and the likelihood of H occurring given X and Y believing H, the same (by substituting the values expressed by (4) and (5) in (3), (3) reduces to (1)) (Goldman 2001: 101).

To go back where we started from, if a layperson has to decide between conflicting experts, according to Goldman, the numbers are only helpful if she has reasons to believe that the two experts are at least partially independent of one another in their judgement. And this, for instance, implies that if they used exactly the same method to arrive at their result, or if they believed the one expert because they considered them epistemic authorities, then their opinions would not be conditionally independent. This would significantly restrict the potential applications of the method of aggregation of additional experts' opinion as a way for laypeople to judge about a domain relative claim: how would a layperson justifiably believe whether experts agreeing with each other have arrived at the same belief independently, or if they used the same methods, or if they believed it because they thought the expert making that claim to be an epistemic authority to them? This information regarding the experts' belief-forming process may not be easily available to laypeople. For instance, how could people know whether all the doctors arrived at their belief that vaccines are efficacious and safe, independently of one another? The method of 'going by the numbers' may not be of much practical use for laypeople deciding what to believe in the case of experts' disagreement, such as the case of vaccines. However, I don't believe that Goldman's analysis is fully correct. Or better: I believe that the partial independence principle, even though correct, leaves room for giving evidential weight to additional official experts' agreement in any case, regardless of whether official experts are non-discriminative reflectors (because they employ the same method to arrive at the belief, or because they treat the first expert as an authority) or not. I will argue for this in what follows.

Despite the intuitive appeal of Goldman's independence principle²⁶, the requirement that the believers who agree on a claim need to be at least partially independent in order for their agreement to have evidential weight in support of the truth of the claim they make, has been criticised by David Coady (2006) and Jennifer Lackey (2013). Let's start with Coady. Coady observes that Goldman's principle is too strict, because it does not consider the possibility that the novice, throughout the argument, changes his confidence

²⁶ Others who have defended a similar principle are Elga (2010) and Kelly (2010).

in the accuracy of the first expert ($P(H)|X(H)$) (2006:71). Coady argues that the layperson's opinion on H could be swayed by Y's agreement, if the 'non-discriminating reflector' Y was a reliable judge of X's expertise, even though Y would have followed X regardless of the truth of H (since Y believed X to be a very good expert). So, Coady concludes, the non-independence principle would be generally true only if one could assume that "followers are invariably unreliable judges of gurus", which we cannot do (2006: 72). A similar point, even though restricted to the case of peer disagreement, has been made by Lackey (2013)²⁷. Lackey argues that numbers can be significant even when the belief-holders are dependent of each other, because the additional support of peers expresses the supporters' belief in the reliability of the source on which they decide to depend. Lackey distinguishes 'autonomous' from 'non-autonomous' complete source dependence, where

The autonomous version of this dependence involves a subject exercising agency in her reliance on a source of information, critically assessing its reliability, monitoring for defeaters, and comparing the content of the belief that she forms with her background beliefs. This, I take it, is the minimum required for rational belief formation. (Lackey 2013:249)

The autonomous decision to rely on a source speaks in favour of the reliability of that source. Lackey argues that this 'autonomy principle' can adjudicate all the cases in a way that is aligned to our intuitions: if an agent autonomously chooses to completely rely on a peer this is evidence for the reliability of the source, if the agent simply parrots the source, then this is irrational and does not provide additional support for the source's claims. The flaw she identifies in Goldman's account is the same that Coady pointed out:

These considerations reveal why it may be misleading to focus on a case such as the guru and his blind followers, as Goldman does. Typically, blind followers are not only non-discriminating when it comes to a particular opinion of their guru's, they are also non-discriminating of their guru's opinions across the board and perhaps even of gurus in general. (Lackey 2013: 262).

Apart from the guru-followers case, Lackey and Coady emphasise that there are many situations in which the independence principle is satisfied, and yet the dependent belief-

²⁷ Lackey restricts her focus to the relevance of numbers for disagreement among peers, but for the purpose of our discussion we can ignore this restriction and just keep focusing on numbers in disagreement in general.

holders provide evidence that the source is reliable: numbers matter more often than Goldman allows for.

As a reply to Coady and Lackey, I first want to notice that Goldman *does* allow for the sort of information that they believe is left out by independence to affect a layperson evaluation of competing experts: in Goldman's terms this is the agreement from 'meta-experts'. Meta-experts are just people who are reliable at identifying experts. If the people agreeing with one expert on a claim do so because they are reliable meta-experts, this kind of information should be taken into consideration, even in Goldman's account. However, they are right to point out that Goldman treats approval by meta-experts as a completely separate question, and that so formulated, the independence principle leaves these cases out. However, I don't believe that the independence principle is false, and needs to be substituted. I want to argue that the independence principle is essentially correct, even though its formulation needs to be slightly refined to accommodate for meta-expertise.

According to Goldman, an opinion holder is totally dependent on another if they are just as likely to endorse their view if the view is true than they are if it is false:

$$P(Y(H) | X(H) \& H) = P(Y(H) | X(H) \& \neg H)$$

However, as Coady and Lackey point out, there may be cases in which this principle holds, and yet the agreement of Y is significant to the truth of H. An obvious case is when the people non-independently agreeing with the source are meta-experts. But this is not the only case: imagine a group of scientists, who all decide to adopt the same method to determine whether a certain hypothesis H is true. After applying the method, each in their laboratory, they conclude that the hypothesis is true. In this case, the principle of non-independence holds: if the method was incorrect and the hypothesis was false they would still agree with each other. However, the fact that a group of scientists decided to employ this method and not another is still epistemically significant: it tells in favour of the reliability of the method, and in turn, of the truth of the hypothesis. Hence, the agreement between the scientists would seem to be evidence that the hypothesis is true. However, I don't believe that the principle of independence delivers an incorrect verdict in these situations: in these cases what happens is that, even though the opinion holders are totally dependent on one another on the claim at issue (H), they *are not* totally dependent on the claim that led them to choose the source on which they are dependent ('M is a good method for determining whether H', 'A is an

epistemic authority for me on claim H'...). The independence principle would have been satisfied *for these claims* since the scientists are at least partially independent of one another in the choice of the method to employ (as a simplification to treat these cases together, relying on another expert can be considered as a 'method' of deciding whether H): they are more likely to agree that the method (or expert) is good if it is true that the method is good than if it is false. Strictly speaking Goldman is right: if non-independent opinion holders agree on a claim, this is not evidence in favour of *that claim*. The fact that all the scientists agree on H after having used the same method to determine whether H, *is not directly evidence for H*. However, if the scientists are at least partially independent of each other in their opinion that M is a good method, their agreement on the reliability of method M is evidence that M is a good method for determining whether H, which in turn is evidence for the truth of H. The same holds for the agreement of the experts in relying on the superior expert. In other words, the independence principle seem to deliver a verdict in these situations that is aligned with our intuitions.

We can incorporate these observations in Goldman's account by claiming that the agreement of Y with X's belief that H is evidence in favour of the truth of H if

- (i) Y is at least partially independent of X in forming the belief that H

$$P(Y(H) | X(H) \& H) > P(Y(H) | X(H) \& \neg H)$$

- (ii) or if Y is at least partially independent of X in forming the belief M that 'm is a reliable method to determine whether H'

$$P(Y(M) | X(M) \& M) > P(Y(M) | X(M) \& \neg M)^{28}$$

In other words, Y's agreement with X is epistemically significant to determine whether H if Y arrives at H at least partially independently of X, or if Y is at least party independent of X in his belief that X's method m is reliable. Even though (ii) is not *directly* evidence for H, it is indirectly so, insofar as M (X's way of assigning credences to

²⁸ It could be that Y holds the belief M completely dependently of X(M), but Y's support would still count in favour of H if Y believed M dependently of X because Y judged that 'X is a judge for M-matters' (call this M1). Again, it could also be that Y believes M1 completely dependently of X, but Y's agreement would still count in favour of H if Y believed M1 dependently of X because Y judged that 'X is a reliable judge of M1-matters' (M2), and so on. Y's agreement counts in favour of H if there is stage in the chain of dependence in which Y is not completely dependent of X in judging a claim M_n that indirectly supports the truth of H (no matter how many steps separate M_n from H).

H-matters is reliable) is evidence in support of the truth of H. The cases identified by Coady and Lackey do not show that Goldman's independence principle is false: it is just applied to a different proposition, which in turn may indirectly support the truth of the claim at stake. Goldman's account can be refined to account for these cases, by adding a condition that includes in the cases in which additional support is evidence for the truth of a claim, the cases in which the support is given in an indirect way in the form of evidence for the reliability of the source.²⁹

When a layperson is caught between experts' disagreement, the agreement of additional experts on one side or the other could count as evidence in favour of believing the experts that has the wider support, provided that condition (i) or (ii) is satisfied for the additional supporters. But how can a layperson acquire evidence that one of the two conditions is satisfied? If evidence of type (ii) was as hard to acquire for a layperson as evidence of type (i) we would not have made much progress, and the method of number aggregation would not be of much practical use to laypeople, not even after this refinement. However, I believe that it is very often possible for laypeople witnessing experts' disagreement to acquire evidence of the kind (i) or (ii). It may not be easy for laypeople to tell which one of the two kinds it is, but this does not matter, since either of them is sufficient for additional support to count as evidence for the truth of the experts' claim. What does it mean for a layperson to acquire evidence of type (i) or (ii)? A layperson would have to learn that the believers of the jointly held propositions have either independently formed their belief, or have dependently done so, but independently formed the belief in the reliability of the method employed to form the belief (including as a method the reliance on another person). How can they obtain such information?

Again it is helpful to look at official experts and the ways in which they conduct scientific research. In the eighth chapter of "The Advancement of Science", Philip Kitcher investigates how scientists should respond to each other, and under what circumstances belief on authority is rational for them (1993:304). Kitcher argues that there are two ways in which experts can assess the credibility of each other to decide under what circumstances to defer to the authority of a fellow expert: *direct calibration* and

²⁹ A similar revision of Goldman's account has been proposed by Barnett (forthcoming). He proposes a principle according to which additional agreement is significant if the two opinion-holders agreeing with each other could not have been expected to agree in advance.

indirect calibration. In direct calibration the experts base the evaluation of another experts' expertise through their own opinion on the topic. In indirect calibration, they base their opinion on other experts whom they have assessed through direct calibration (1993:316-322). Given that official experts are generally reliable in their fields (at least in the scientific fields we are considering, such as medicine, in which the advancement and technological innovations in the field are evidence of the general reliability of their experts), using their own opinion directly or indirectly to evaluate other scientists should bear reliable results. Moreover, the official experts' connectedness in a field makes them able to employ the indirect methods of expertise assessment which we discussed in the previous section, and that have only limited applications for laypeople. Official experts can often witness other experts' debating and can understand their arguments, they can evaluate the evidence of past track records, and their interests and biases. When official experts agree, this is generally evidence that either condition (i) or (ii) is satisfied: either they formed their opinion independently, or they totally relied on a source that they considered reliable. Official experts are generally reliable meta-experts. Hence, official experts' agreement is generally (unless there is reason to think otherwise) epistemically significant to determine whether a certain claim is likely to be true. We can conclude that, in the case of experts' disagreement, 'going by the numbers' seems to be the most reliable and accessible way for laypeople to decide what to believe. Once appropriately refined, the independence principle is such that we have *prima facie* reasons to believe it holds for all official experts. Evidence of consensus of official experts is easier to acquire than any other kind of evidence we have considered so far. In the case of vaccinations, for instance, it is clear from the policy choices made by different countries that the majority of official experts community-wide believes vaccines to be effective and safe. Thus, the rational thing to do for laypeople is to 'go by the numbers' and trust the majority of official experts.

Concluding, laypeople have ways to assess the credibility of rival experts without engaging with the content of their arguments (which would require them to become experts in the field). I have argued that the easiest and most reliable way to assess rival experts is considering the additional agreement of other official experts. Such agreement is either evidence in support of the proposition at stake, if the experts have arrived at their belief independently of each other, or is evidence in support of the superior expertise of one side of the dispute, given that official experts are generally reliable meta-experts. By refining the principle of independence I have proposed a way to put

together Goldman's two methods of (B) agreement by other experts and (C) appraisal by meta-experts. Considering the agreement of official experts is not the only way in which laypeople can employ methods (B) and (C). However, given the lack of expertise that characterises laypeople, it may not be easy for them to acquire evidence that people other than official experts are reliable meta-experts or have arrived at the jointly held belief at least partially independently of each other. Looking at the agreement of official experts looks like the most promising way for laypeople to decide in cases of disagreement. Where does this leave our analysis? After having outlined my account of epistemic authority, I have considered possible ways that may cast doubts upon the authoritativeness of an official expert on a specific claim. I have considered specific doubts regarding one of the conditions of authority not being fulfilled for a claim or set of claims, and argued that in these cases it is rational for laypeople to consult another official expert (or someone for whom they have reasons to believe that the conditions of epistemic authority hold). I have then considered the case of disagreement among experts who individually taken satisfy the condition of epistemic authorities for a layperson. I have discussed Goldman's (2001) methods for assessing expertise, and argued that the most promising method for laypeople trying to assess competing authorities seems to be relying on the side which has most supporters among official experts. For vaccinations, the situation seems clear: vaccines are deemed safe and efficacious by the majority of official experts, and this is what laypeople should believe too.

There is still one source of scepticism in official experts that I would like to address in this paper. In section 3, I argued that the community of official experts provides laypeople with coverage for domain related propositions, and hence that in the absence of official experts' reports on a claim, it is never permissible for laypeople to assign positive credence to it. In the next section, I want to consider the case in which laypeople may be sceptical that the conditions of coverage reliability are met.

6. Doubts about official experts' epistemic coverage: when is a rumour reliable?

In section 3, I have argued that, for official expert, we have reasons to believe that the following conditions are satisfied:

- (i) The community of official experts in domain D reliably tracks p-relevant evidence and reliably determines whether p on that basis. (Where p is any question

within the domain that is deemed worthy of being inquired by official experts in that domain)

- (ii) The community of experts reliably makes any results of their investigations available if it meets the standards of good scientific practice (by means of peer-reviewed publication).
- (iii) Were discoveries to be made regarding p , they would be made available by the official experts by the time laypeople relied on them.

I have claimed that these conditions, jointly considered, ground the idea that it is never permissible for a layperson to assign a positive credence to a claim, if no official report in favour of that claim exists. Moreover, if the layperson has reason to believe that an investigation has happened regarding the matter, and sufficient time passes between the beginning of the investigation and the time when the layperson registers the absence of experts' reports on the matter, this is enough to ground the layperson's absence-based belief that the claim is false.

In this section, I want to consider a controversial situation: the situation in which, even though no official reports make a certain claim, laypeople encounter *rumours* that claim the truth of it. If the rumours are particularly widespread, a layperson may wonder why so many people believe them, and whether there may be some truth to them. Moreover, if the rumour is judged likely to be true by many seemingly reliable people, this can be taken by the layperson as evidence that the conditions of epistemic coverage may not be met by the experts' community for that precise claim. One such example is the vaccines-autism controversy. Despite the fact that, to date, no certain link between autism and

vaccines has been claimed by official experts³⁰, many non-official reports circulate on the topic, and many people wonder about their truth. The tenaciousness and plausibility of these rumours makes some people wonder if it may be the case that scientists are not investigating this topic enough, or if they may be acting out of interest in not making the results of their studies public. Despite the reputation of rumour as unreliable, it has recently been argued by some social epistemologists (Coady 2006; 2012 ch.4; Dentith 2013) that rumour is a much more reliable way of forming beliefs than previously thought, and that ‘the fact that a proposition is rumoured to be true is evidence in favour of it being true’ (Coady 2001:42). If they are right, the persistence of rumours about the existence of facts that are not confirmed by official experts could constitute a serious source of doubt in the coverage reliability of the community of experts, and should thus be taken seriously. But how can laypeople assess the reliability of rumours, without being able to confirm or disconfirm them directly, due to their lack of expertise?

Coady describes rumour as a kind of communication characterised by two essential features: (i) rumour is communication that has been passed on by many people and (ii) it

³⁰ This needs some clarification. The first paper to suggest a link between vaccines and autism is Wakefield et al. (1998). The paper was later retracted from the journal on the basis that it made use of falsified data (see Deer 2011), and Andrew Wakefield removed from the UK medical register in 2010. So, strictly speaking, Wakefield’s paper does not count as an official report, nor is he an official expert anymore. A few other published papers exist suggesting the relation between vaccines and autism, by Shaw and Tomljenovic. Shaw and Tomljenovic have a history of retracted (and sometimes republished) papers (2014; 2016; 2017). Two of their papers (2011a, 2011b) were discredited by the WHO which found them ‘seriously flawed’ (see http://www.who.int/vaccine_safety/committee/topics/adjuvants/Jun_2012/en/). Even though this is a controversial issue, I believe that we can assume for sake of the argument that no official report exists that claims that vaccines cause autism. In fact, even the existing ones do not claim this as bluntly, and only connect the presence of aluminium in the brain— which is contained in most processed foods and cosmetics in addition to being used as an adjuvant in some vaccines— with autism. For the sake of my argument we can ignore these complications and simply assume that there is no such official report. If this assumption sounds unwarranted we can restrict my example to the MMR controversy, or the Thimerosal one. There it is clear that there is no official report (after Wakefield’s discarded one) claiming that the MMR vaccine or Thimerosal causes autism— and some people still believe that they do. The reason why we need to assume that no official report exists is that this is what is necessary for a claim to be a rumour. If official reports exist, then it is not a case of rumour, but a case of disagreement, and it is sufficient to look at the additional support on each side.

has an unofficial status (2006: 49)³¹. Nothing in this definition implies that rumours are necessarily unreliable. In fact, according to Coady, the spreading of rumours can be seen as ‘filtering’ mechanism that should count as evidence of their truth:

A would-be rumour-monger can be more or less reliable. All else being equal, the less reliable he is, the less likely it is that he will be believed, and the less likely it is that the rumour will be passed on. Furthermore, if it is passed on, it is likely it will be passed on to few people and with a low estimate of plausibility. Hence, all else being equal, the greater the reliability of those who spread a rumour, the more likely it is to survive and spread. Hence, if you hear a rumour, it is not only *prima facie* evidence that it has been thought plausible by a large number of people, it is also *prima facie* evidence that it has been thought plausible by a large number of reliable people. And that really is *prima facie* evidence that it is true. (Coady 2006: 47)

If a person decides to pass on a rumour, this should be considered *prima facie* evidence that the source from which she heard the rumour is considered by her reliable³². The first thing to notice about this argument is that even if Coady was right that people are more likely to pass on rumours if they believe the source to be reliable (which is itself a questionable claim), this only counts in favour of the truth of a rumour if we can also assume that these people are reliable at assessing the reliability of other people. Even granting the epistemic integrity and vigilance of the rumour-mongers, the fact that a rumour has spread is only *prima facie* evidence that it has been thought plausible by a large number of people who believed that the speaker from which they heard the rumour was reliable. However, this is *not* evidence that the rumour itself is true, since it may as well be that people are not reliable at estimating the reliability of the speakers from which they hear the rumour. For rumour in general, it is not so straightforward as Coady assumes that it more likely for true rumours rather than false ones to survive.

But even if his argument was correct, it would still be different for the case of rumours regarding topics that fall within highly specialised domains, in which expertise is required to form reliable opinions (such as the field of vaccines). We need to distinguish two sense of reliable: a rumour-monger can be reliable in the sense of being an *honest testifier* who only spreads rumours that she genuinely believes to be plausible; and in the

³¹ Similarly, Dentith defines rumour as ‘An unverified proposition which has been heard by an agent and then expressed to another agent’ (2013: 2)

³² A very similar argument is made by Dentith (2013).

sense of being *accurate*, i.e., she spreads mostly true information. Coady's argument for the reliability of rumour uses reliability in the former sense. What does it mean for a claim to be plausible to a hearer? A plausible claim is one that does not contradict the hearer's other beliefs (see Dentith 2013: 2). In domains which do not require expertise (such as office gossiping), the fact that people spread rumours that they find plausible and that they heard from reliable (honest) testifiers, could be evidence of the truth of the rumour³³, because in these fields plausibility bears some sort of relation to accuracy (or at least Coady and Dentith seem to believe so). However, in highly specialised domains of expertise, plausibility does not imply accuracy, and the fact that *a non-expert* finds a claim plausible, is not evidence in support of that claim. This is because in these domains non-experts are not reliable opinion-holders, and have little background knowledge against which to test the plausibility of claims. So, either the rumour mongers are experts, in which case the judgement of plausibility may support the truth of the claim, or what they are doing is simply passing on something they found plausible and they heard from an honest testifier, who also found the claim plausible (and so on). Since there is no reason to assume that all the people involved in spreading the rumour are (non-official) experts, even if Coady was right in his argument, it would still bear no significance for domains which require expertise for judgements of plausibility to be significant. Moreover, for rumours in non specialised domains (like office gossiping) the interaction among the rumour-mongers can improve the original rumour: people can put together partial information they have, they can make sense of information gaps, and so on³⁴. However, this is not the case for rumours in domains of expertise: since these fields require expertise, non-experts trying to make sense together of partial information they have and of gaps in their information, will most likely produce inaccurate results.

Since no filtering effect can be produced by rumouring in specialised fields of expertise, and the fact that a rumour spread is only evidence that many people found it plausible and heard it from testifiers they thought were honest, the number of people who believe the rumour, and the fact that it keeps spreading is *not evidence* of its truth. Other factors can explain why the rumour spreads that do not imply its relation to the truth. In his book 'On Rumor' Cass Sunstein analyses different factors that explain why

³³ It crucially depends on whether plausibility is a guide to truth, but I shall not engage with this part of the argument here, and just assume that Coady is right.

³⁴ See Shibutani (1966).

rumours— both true and false— propagate. He claims that sometimes rumours spread as *informational cascades*, in which “once a certain number of people appear to believe a rumor, others will believe it too, unless they have good reason to believe that it is false.” (2008:30). Since rumours are often about topics for which individuals lack direct personal information, they defer to the crowd. The larger the crowd the more people defer to it, thinking that ‘so many people cannot be wrong’. And this is how an informational cascade is generated, in which true and false claims alike are propagated (Sunstein 2008: 33). Another way in which rumour spread according to Sunstein is through *conformity cascades* in which people decide to propagate a rumour even if they don’t believe it themselves, just to act in conformity with a group. Conformity cascades can help false rumours to spread, especially in closely knit communities (Sunstein 2008:39). Another factor that facilitates rumours spread is group-polarisation, in which the conviction in a rumour is strengthened by discussion with like-minded people (Sunstein 2008:45). All of these mechanisms explain how false rumours can spread just like true ones. Social and emotional factors, together with biases and motivation, can explain why a rumour is still around, and keeps spreading, without being true.

For now I have only considered rumours whose content falls within a domain of expertise, and argued that these rumours can be explained in ways that do not imply their truth-likeness. However, laypeople may be sceptical of the epistemic coverage of the experts because of another type of rumours: rumours regarding the experts community’s honesty. For instance, in the case of autism and vaccines, laypeople may wonder whether the rumour according to which all the official experts are corrupted by the big pharmaceutical companies is evidence for the truth of this claim. I believe that two factors count against the reliability of rumours regarding the community of experts’ honesty: first, there is the fact that laypeople are not only poor judges of the claims made by experts in their fields of expertise, but also of the experts *themselves*, and of their interests and biases. Laypeople often have limited knowledge of how scientific research works, how it is funded, and what interests are at stake in a specific experts’ research. Hence, also for experts’ honesty and integrity laypeople’s judgements of plausibility do not bear epistemic significance. The rumour that all official experts are corrupted is at best evidence that many people who heard the rumour from someone who they thought was honestly expressing their opinion on the matter, thought it was plausible that all the experts are corrupted by the big pharmaceutical companies. And this does not imply the truth of the claim, unless there is evidence that the judgement of plausibility of

laypeople on experts' interests have a relation to the truth. Secondly, in this case too the survival of the rumour seems best explained by appealing to the factors mentioned by Sunstein (2008), especially if we consider the environments in which these rumours spread, which are usually online forums of closed groups of like-minded people.

In this section I have considered rumour as a potential source of scepticism in the epistemic coverage provided by the experts' community. I have argued that for rumours whose content falls within domains where expertise is required for reliability, the survival of a rumour is not evidence of its truth. Moreover, I argued that this is often the case also for rumours regarding the experts' honesty, since the background information of laypeople regarding the interests and biases of scientists is very limited. The survival of these kinds of rumours in our societies are best explained by social and affective factors such as those identified by Sunstein (2008). However, rumours are not the only thing that may lead laypeople to be sceptical of the epistemic coverage provided by the official experts. Another (possibly the most widespread) way in which epistemic coverage can be doubted, is by appealing to a conspiracy explanation.³⁵ In a conspiracy explanation, a person explains a phenomenon by appealing to the existence of a conspiracy. For instance, a conspiracy explanation would be the theory that vaccines cause autism, and the reason why the community of experts doesn't say anything about this, or even denies it, is that they are corrupted by the big pharmaceutical companies, who have interests in people believing that vaccines are safe. The difference between a theory like this, and the rumour we have considered in the previous paragraph, is that the rumour would be accepted by the believers on *testimonial grounds*, on the basis that they heard it from a trusted source, found it plausible and they believed that the fact that so many other people found it plausible must be evidence of its truth. Differently, a conspiracy explanation is believed because is taken to be *best explanation for the evidence available* by its believers³⁶. A person can of course learn the conspiracy explanation from a source, in a testimonial way; however, the *reason* for believing the conspiracy explanation is not testimonial, differently from the reason to believe a rumour. So, in the case of the vaccines-autism controversy, the rumour that the scientists are corrupted would be

³⁵ I want to avoid the term conspiracy theory, since the term is ordinarily associated with irrationality, and I don't want to assume that these theories positing a lack of epistemic coverage due to a conspiracy are necessarily irrational.

³⁶ This distinction is proposed by Dentith (2013: 18-19). He uses this distinction to argue that rumours, differently from conspiracy theories, are a way of fact-checking.

believed on the basis that it is found plausible and that so many other people believe it; while the conspiracy explanation would be believed on the grounds that it is taken to be the best explanation for the evidence available.

Conspiracy theorising is too broad a topic to be fully investigated in this paper, and I shall not provide a definitive solution to this problem here. However, let's just focus on an element of the conspiracy that is supposed to establish the corruptness of all the official experts of the world. Allegedly, this theory is proposed by its supporters as an *inference to the best explanation* from all the evidence available to the lack of epistemic coverage: the argument is that all the evidence available is best explained by appealing to a conspiracy according to which the community of official experts has been corrupted by the pharmaceutical companies to deny a relation between vaccines and autism³⁷, to the effect that the community of official experts does not provide epistemic coverage for laypeople on this topic. So, the proposers of this explanation should not deny the official experts' epistemic coverage independently of the inference they make. In fact, the mere evidence that pharmaceutical companies have economic interests in vaccines is not alone enough to support the claim that all the official experts of the world are corrupted—I believe nobody would think that it is, not even the supporters of this theory. So, for an inference to the best explanation to be valid, it needs to be the case that some facts are left unexplained by the official theory (abductive trigger), and that the new theory (in this case, the conspiracy) can explain all of the facts in a better way than the rival explanation (in this case, the official story, that vaccines do not cause autism).³⁸ What are the facts that are supposedly left unexplained by the official story? From what I could learn on the internet, these facts include: the reports of parents whose children were diagnosed with autism after having been vaccinated, the fact that autism occurrence has increased in the population over the past twenty years, the fact that the few experts who claimed such a relation have been harshly criticised³⁹. Now these facts *do* have an explanation in the official theory. I believe that you don't need to know what an inference to the best explanation is to understand why the conspiracy hypothesis is not the best explanation

³⁷ I stick to the example of vaccines and autism, but this is not the only conspiracy around vaccines. A similar reasoning can be applied to the other conspiracies to check whether those are more plausible than this one.

³⁸ For a more detailed discussion of inference to the best explanation, see for instance Lipton (2007:239-240)

³⁹ I take these to be the main reasons, but it may as well be that there are others.

for these facts, given how many people would need to be corrupted in order for this conspiracy to be happening undisturbed, and given how scientific research and funding works— and given that these facts are already explained in the official account. Moreover, the nature of the fact mentioned as ‘abductive triggers’ makes it such that *expertise* is required to judge what the best explanation is for them, and given that I don’t *already* have evidence that the official experts are corrupted, I should trust the official experts’ explanation. I believe that the conspiracy explanation is only *really* convincing if you *already are* sceptical of the fact that the community of official experts and the media provide epistemic coverage for this and similar issues. And if you already are sceptical of this, then my analysis of epistemic dependence on experts will not have much significance for you.

Conclusions

Given the complexity that characterises the human knowledge in our times, we often encounter topics that require that we rely on experts to form an opinion. But how exactly are we to depend on experts in these situations? In this paper I have tried to provide an account of the relation between experts and laypeople, on the basis of both epistemological considerations, and consideration about the way in which the cognitive labour is actually divided in our societies. Through the notion of official expertise, I have outlined an account of epistemic authority that is grounded in the way in which laypeople actually encounter expertise in their lives. I have argued that laypeople have *prima facie* reasons to treat official experts as epistemic authorities in all of the questions within their domain of expertise. In the absence of reasons to think otherwise, laypeople should preemptively adopt the official experts’ beliefs, simply on the grounds that they are official experts. I have then considered different reasons for laypeople to be sceptical of experts’ opinions, and the ways in which laypeople can decide what to believe in these situations, without having to engage with the content of the experts’ claims. The result for vaccines scepticism, in a nutshell, has been that laypeople have reasons to trust the majority of the official experts, and mistrust any rumours regarding alleged facts about vaccines that are not confirmed by official experts.

There are a few issues that I would like to address in concluding this paper. First, this paper relied on an optimistic assumption regarding the reliability and honesty of official experts and the efficacy of the institutional monitoring processes for experts’ reliability

(such as peer review). This implies that if the reader does not share this intuition, this work will be of little significance to the way they think they should conduct their epistemic lives. However, my assumption is not completely arbitrary. As I have mentioned already a few times throughout this work, the domains to which this discussion applies are scientific domains, whose achievements can be directly or indirectly (though their applications to other fields) assessed by laypeople, either through the evaluation of technological advancements or through the predictions they make about observable facts. I take this to *prima facie* grant some form of trust in the reliability of the community of official experts in these domains. This does not mean that interests never play a role in these domains and that experts are never corrupted. These cases are covered by my analysis. However, once the optimistic assumption is granted, specific reasons are necessary to mistrust specific experts and their claims. Moreover, I have not assumed that official experts are always reliable: given the lack of expertise that characterises laypeople, they will be better off deferring to the experts even if the experts sometimes make false statements. Even though further research would be necessary for me to establish that official experts are generally reliable, it seemed to me a reasonable assumption to make.

Secondly, it may be argued that this account of epistemic dependence would lead to epistemically catastrophic results if it were applied in societies where research and expertise are organised in less favourable ways. Again, I want to stress that the fact that I have relied on how the production and spread of knowledge is organised in *our* societies restricts the scope of this paper significantly. For instance, this is no indication of how the relation between experts and laypeople should be shaped in, say, totalitarian states with censorship. This account of dependence on experts would be totally unjustified in different situations. Thirdly, I want to make clear that this is exclusively an account of the dependence of *laypeople* on experts. I have entirely left out considerations about how experts should respond to each others' opinions. While it may at first glance appear 'conservative', this account does not claim that experts should defer to each other and to the majority in the same way.

Finally, I want to consider the worry that my account does not have a place for the value of laypeople's understanding of the issues, and for their autonomous decision and critical thinking. In this respect, there are two things I want to notice. First, I want to make clear that I stand by this claim, and I am willing to take it even further. Engaging

with the arguments made by experts, attempting to understand the issues at stake, and finally making up your own mind as a critical thinker is not a good strategy to achieve reliable beliefs in specialised domains of expertise. It does not matter for epistemic purposes that laypeople understand the claims they hold and the reasons why they are the right ones. Very often, the impression of understanding that laypeople obtain by engaging with the domain-relative material makes the expertise gap between them and the experts look much smaller to them than it is in reality. And this cannot be beneficial for them for forming correct beliefs in specialised domains, nor for the experts who have to spend time trying to make laypeople understand their mistakes. In other words, my impression (which surely would require some empirical research to be confirmed) is that laypeople's understanding of the issues is not just neutral to their acquisition of accurate beliefs in specialised domains, but at least in many cases, *detrimental* to it. I believe that a much better strategy would be for laypeople to invest that cognitive effort in picking the right experts to defer to. However, the fact that understanding and critical thinking have no place in an *epistemic* account of the dependence on experts does not necessarily exclude these values altogether from the relation between experts and laypeople. It only shows that our epistemic concerns and our other values may come apart: we value autonomy and understanding, together with truth and rationality, and sometimes we have reasons to opt for a suboptimal epistemic conduct in favour of other things we value.

At the same time, these considerations point towards possible developments of this line of research to identify ways in which institutions could address controversies such as the vaccinations case. If my account is correct in claiming that the rational way to respond to official experts is to defer to their authority, and if the rationality of this strategy could be recognised by laypeople, then this would seem to suggest that a good strategy for addressing those cases in which laypeople's opinion in specialised domains differs from that of the majority of experts, would be to work on improving laypeople's perception of official experts reliability, and of the ways in which institutions monitor experts' honesty and conflict of interests, rather than to present laypeople with more domain-relative facts and information.

Bibliography

Almassi, Ben (2012). “Climate Change and the Ethics of Individual Emissions”. *Perspectives* 4 (1):4-21.

Allport, Gordon W., & Postman, Leo (1947). *The Psychology of Rumour*. New York: Henry Holt & Co.

Amin, A. B., Bednarczyk, R. A., Ray, C. E., Melchiori K. J., , Graham, J., Huntsinger, J. & Omer, S. B. 2017. “Association Of Moral Values With Vaccine Hesitancy”. *Nature Human Behaviour* 1: 873–880

Barnett, Zach (Forthcoming). “Belief Dependence: How Do the Numbers Count?” *Philosophical Studies*:1-23.

Bergmann, Michael (2005). “Defeaters and Higher-Level Requirements”. *Philosophical Quarterly* 55 (220):419–436.

Coady, David (2006). “When Experts Disagree”. *Episteme* 3 (1-2):68-79.

Coady, David (2006). “Rumour Has It”. *International Journal of Applied Philosophy* 20 (1): 41-53.

Coady, David (2012). *What To Believe Now: Applying Epistemology to Contemporary Issues*. Wiley-Blackwell.

Coady, Cecil Anthony John (C.A.J.) (2006). “Pathologies of Testimony.” In Jennifer Lackey & Ernest Sosa (Eds.), *The Epistemology of Testimony*. Clarendon Press.

Deer, Brian (2011). "How the Case Against the MMR Vaccine Was Fixed". *BMJ* 342: C5347.

Dentith, Matthew (2013). "Have You Heard? The Rumour as Reliable". In Greg Dalziel (Ed.), *Rumour and Communication in Asia in the Internet Age*. Routledge. Pp. 46-61.

Dormandy, Katherine (Forthcoming). "Epistemic Authority: Preemption or Proper Basing?" *Erkenntnis*:1-19.

Elga, Adam (2007). "Reflection and Disagreement". *Noûs* 41 (3):478–502.

List.

Elga, Adam (2010). "How To Disagree About How To Disagree." In Ted Warfield & Richard Feldman (Eds.), *Disagreement*. Oxford, UK: Oxford University Press. pp. 175-186.

Feldman, Richard (2006). "Epistemological Puzzles About Disagreement." In Stephen Hetherington (Ed.), *Epistemology Futures*. Oxford University Press. pp. 216-236.

Gelfert, Axel (2013). "Coverage-Reliability, Epistemic Dependence, and the Problem of Rumor-Based Belief". *Philosophia* 41 (3):763-786.

Goldberg, Sanford C. (2010a). *Relying on Others: an Essay in Epistemology*. Oxford University Press.

Goldberg, Sanford C. (2010b). "The Epistemology of Silence". In Adrian Haddock, Alan Millar & Duncan Pritchard (Eds.), *Social Epistemology*. Oxford University Press. pp. 243–261.

Goldberg, Sanford C. (2011). "If That Were True I Would Have Heard About It by Now". In A. Goldman & D. Whitcomb (Eds.), *The Oxford Handbook of Social Epistemology*. Oxford University Press.

Goldman, Alvin I. (1999). *Knowledge in a Social World*. Oxford University Press.

Goldman, Alvin I. (2001). 'Experts: Which Ones Should You Trust?' *Philosophy and Phenomenological Research* 63 (1):85-110.

Grundmann, Thomas (Manuscript). "Epistemic Authority: Preemption Through Undercutting Defeat."

Guerrero, A., 2016. "Living With Ignorance in a World of Experts". In Peels, R. (Eds.) *Perspectives on Ignorance From Moral and Social Philosophy*. Routledge.

Hardwig, John (1985). "Epistemic Dependence". *Journal of Philosophy* 82 (7):335-349.

Hardwig, John (1991). 'The Role of Trust in Knowledge'. *Journal of Philosophy* 88 (12): 693-708.

Hazlett, Allan (2016). "The Social Value of Non-Deferential Belief". *Australasian Journal of Philosophy* 94 (1):131-151.

Hsu, Anne S.; Horng, Andy; Griffiths, Thomas L. & Chater, Nick (2017). "When Absence of Evidence Is Evidence of Absence: Rational Inferences From Absent Data." *Cognitive Science*.

Jäger, Christoph (2016). Epistemic Authority, Preemptive Reasons, and Understanding. *Episteme* 13 (2):167-185.

Joyce, James (2007). "Epistemic Deference: the Case of Chance". *Proceedings of the Aristotelian Society* 107 (2):187 - 206.

Kelly, Thomas (2010). "Peer Disagreement and Higher Order Evidence." In Alvin I. Goldman & Dennis Whitcomb (Eds.), *Social Epistemology: Essential Readings*. Oxford University Press. pp. 183--217.

Keren, Arnon (2007). "Epistemic Authority, Testimony and the Transmission of Knowledge". *Episteme* 4 (3):368-381.

Keren, Arnon (2014a). "Trust and Belief: a Preemptive Reasons Account." *Synthese* 191 (12):2593-2615.

Keren, Arnon (2014b). "Zagzebski on Authority and Preemption in the Domain of Belief". *European Journal for Philosophy of Religion* 6 (4):61-76.

Kitcher, Philip (1993). *The Advancement of Science: Science Without Legend, Objectivity Without Illusions*. Oxford University Press.

Kitcher, P. (2010). "The Climate Change Debates." *Science*, 328, 1230–1234.

Lackey, Jennifer (2013). "Disagreement and Belief Dependence: Why Numbers Matter." In David Phiroze Christensen & Jennifer Lackey (Eds.), *The Epistemology of Disagreement: New Essays*. Oxford, UK: Oxford University Press. pp. 243-268.

Lackey, Jennifer (2018). "Experts and Peer Disagreement". In Matthew A. Benton, John Hawthorne & Dani Rabinowitz (Eds.), *Knowledge, Belief, and God: New Insights in Religious Epistemology*. Oxford: Oxford University Press.

Larson, H.J., de Figueiredo, a., Xiahong, Z., Et Al., 2016. “The State of Vaccine Confidence 2016: Global Insights Through a 67-Country Survey”. *EBioMedicine* 12, 295–301.

Lipton, P. (2007). Alien Abduction: Inference to the Best Explanation and the Management of Testimony. *Episteme*, 4(3), 238-251.

Martini, Carlo (2014). “Experts in Science: a View From the ‘Trenches.” *Synthese* 191 (1): 3-15.

Mercier, Hugo (2011). “When Experts Argue: Explaining the Best and the Worst of Reasoning.” *Argumentation* 25 (3):313-327.

Miller, Boaz (2013). “When Is Consensus Knowledge Based? Distinguishing Shared Knowledge From Mere Agreement”. *Synthese* 190 (7):1293-1316.

Millgram, Elijah (2015). *The Great Endarkenment: Philosophy for an Age of Hyperspecialization*. Oxford University Press USA.

Nguyen, C. Thi (Forthcoming). “Cognitive Islands and Runaway Echo Chambers: Problems for Epistemic Dependence on Experts”. *Synthese*.

Pedersen, Nikolaj Jang Lee Linding & Kallestrup, Jesper (2013). “The Epistemology of Absence-Based Inference”. *Synthese* 190 (13):2573-2593.

Pollock, John (1987). “Defeasible Reasoning”. *Cognitive Science* 11 (4):481-518.

Raz, Joseph (1988). “The Morality of Freedom”. *Ethics* 98 (4):850-852.

Shibutani, Tamotsu (1966). *Improvised News: a Sociological Study of Rumor*. Indianapolis: Bob Merrill.

Sunstein, Cass R (2008) *On Rumors: How Falsehoods Spread, Why We Believe Them, and What Can Be Done*. Princeton University Press

Sunstein, Cass R. & Vermeule, Adrian (2009). “Conspiracy Theories: Causes and Cures”. *Journal of Political Philosophy* 17 (2):202-227.

Trout, J. (2004). “Paternalism and Cognitive Bias”. *Law and Philosophy* 24 (4):393-434.

Tomljenovic Lucja, Shaw Christopher A. (2011a) “Do Aluminum Vaccine Adjuvants Contribute to the Rising Prevalence of Autism?” *Journal of Inorganic Biochemistry*, 105: 1489–1499

Tomljenovic Lucja, Shaw Christopher A. (2011b) “Aluminum Vaccine Adjuvants: Are They Safe?” *Current Medicinal Chemistry*, 18(17):2630–2637.

Tomljenovic Lucja, Shaw Christopher A. (2017) RETRACTED: “Subcutaneous Injections of Aluminum at Vaccine Adjuvant Levels Activate Innate Immune Genes in Mouse Brain That Are Homologous With Biomarkers of Autism”. *Journal of Inorganic Biochemistry*.

Tuzet, Giovanni (2015). “On the Absence of Evidence.” in In Christian Dahlman & Thomas Bustamante (eds.), *Argument Types and Fallacies in Legal Argumentation*. Springer Verlag.

Wakefield, Andrew (1998) et al. RETRACTED: "Ileal-Lymphoid-Nodular Hyperplasia, Non-Specific Colitis, and Pervasive Developmental Disorder in Children". *Lancet*. 1998; 351(9103): 637-641.

Zagzebski, Linda (2012). *Epistemic Authority: a Theory of Trust, Authority, and Autonomy in Belief*. Oup Usa.

Zagzebski, Linda (2013). "A Defense of Epistemic Authority". *Res Philosophica* 90 (2): 293-306.

Zagzebski, Linda (2016). "Replies To Christoph Jäger And Elizabeth Fricker". *Episteme* 13 (2):187-194.