

MASTERARBEIT | MASTER'S THESIS

Titel | Title

How could laypeople assess the trustworthiness of personal-
finance advice?
An interdisciplinary study

verfasst von | submitted by
Tobias Derschan BA

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of
Master of Arts (MA)

Wien | Vienna, 2024

Studienkennzahl lt. Studienblatt | Degree
programme code as it appears on the
student record sheet:

UA 066 642

Studienrichtung lt. Studienblatt | Degree
programme as it appears on the student
record sheet:

Masterstudium Philosophy and Economics

Betreut von | Supervisor:

Univ.-Prof. Dr. Dr. h.c. Martin Kusch

Mitbetreut von | Co-Supervisor:

Lic. Dr. Justin Leduc master

How could laypeople assess the trustworthiness of personal-finance advice? An interdisciplinary study.

Contents

Introduction	3
Main part.....	9
1. A conceptual analysis of expertise	9
2. Laypeople and their epistemic justification to trust experts	17
2.1 Who are the laypeople?	17
2.2 Rational trust in epistemic authorities by employing fallible heuristics	19
3. The general discussion of the debate on heuristics for assessing experts	26
3.1 Assessing biases and incentive structures	26
3.1.1 A conceptual analysis of biases and incentives	27
3.1.2 Strategic information transmission via ‘cheap talk’	32
3.1.3 A game theory informed analysis of incentive structures	36
3.2 Directly assessing the past track record of putative experts	41
3.3 Comparing testimonies of multiple experts	44
3.3.1 The epistemic value of agreement and consensus	44
3.3.2 The robustness of expert testimonies.....	47
3.3.3 Indicators for robust testimonial evidence.....	52
3.3.4 Strategically utilizing multiple testifiers to disincentivise dishonesty	55
4. An interdisciplinary apprehension of personal-finance expertise	58
4.1 The discipline of personal finance	60
4.2 Personal finance as a social science	64
4.3 Economic forecasting in personal finance	70
5. Applying the general debate to the domain of personal finance	73
5.1 Assessing incentive structures in personal finance.....	73
5.1.1 The ambiguity and uncertainty of personal-finance advice.....	74

5.1.2 Common incentives in personal finance.....	75
5.1.3 Opportunities to bias personal-finance advice.....	81
5.1.4 A further heuristic utilizing strategic thinking	86
5.2 Assessing the track record of putative personal-finance experts.....	86
5.3 Comparing the testimonies of multiple experts in personal finance	89
5.3.1 Assessing agreement between personal-finance experts	91
5.3.2 Strategically utilizing competition between personal-finance experts	94
Conclusion.....	98
Bibliography	103
Appendix	109

Introduction

Imagine that more money is accumulating in your bank account for the first time in your life than you need anytime soon. What should you do with it? Or what if an old colleague from school is contacting you, whom you have not heard from for 15 years, offering you an exciting investment opportunity? Should you take it? Perhaps you already invested in several such opportunities and are now indebted to your house bank and half your peer group. What would be the most effective way to amortise your debts?

One way or another, many people are confronted with difficult questions on how to manage their money, secure their financial well-being or save for future projects who have no reliable knowledge on these matters. Confronted with inscrutably manifold investment products, insurance policies, and spurious finance tips, being left at one's own devices seems risky. After all, making mistakes with as delicate affairs as one's monetary funds can be literally costly. One straightforward approach to address this issue is to become an expert in personal finance, which is a discipline that concerns itself with the financial well-being of individuals. This suggestion, however, might be as straightforward as it is cynical. Learning an expertise requires extensive amounts of time, energy and money. Moreover, personal finance is not the only issue everyday people are confronted with where they could profit from expert knowledge. Since it is unfeasible to become an expert in everything important, most people are laypeople in most contexts. Following the idea of an (epistemic) division of labour, could laypeople not simply defer to experts for advice? The problem is how laypeople could know who such experts are when they lack any special knowledge. Even if these putative experts are not frauds, competent agents can have interests which conflict with their clients as well.

Are there, then, ways to rationally trust expert testimony which do not amount to blind trust? Scholars have repeatedly addressed similar issues in the history of philosophy. Plato, for instance, discussed the problem of differentiating a real doctor from a charlatan in his Socratic dialogues, especially *Charmides*. Later, Sextus Empiricus, stated the, in his opinion, impossibility of judging skills without being an expert oneself. In the 90s, as philosophers began to focus on the role of learning from testimony again, C.A.J. Coady and Scott Brewer discussed the use of expert knowledge in the context of legal trials. It was, however, Alvin Goldman who in 2001 reinvigorated the debate in its modern form until today. He argued that while laypeople might not be able to directly assess the trustworthiness of expert testimony, they can use heuristics to fallibly gauge which putative experts are more likely to be competent and honest. After all, learning about the fact that an expert is being paid by a specific interest group or has

never engaged with the topic in question before, ought to intuitively diminish her trustworthiness. Indeed, most contemporary participants in the debate agree that laypeople can have better or worse reasons to trust experts and, accordingly, the debate concerns itself with discussing which heuristics are epistemically valuable, to what degree and in which context. Respectively, Goldman called this the novice/ 2 expert problem in contrast to the expert/ expert problem, in which Phillip Kitcher analyses the foundation of trust between scientists. While multiple notions are being used to address this topic, I will from now on simply call it the expert-identification question. (see C. A. J. Coady 1994; Bekker 1842 cited after Goodwin 2011, 288; Brewer 1997; A. I. Goldman 2001; Hardwig 1985; Kitcher 1995; LaBarge 1997; Schleiermacher 1857, Charm. 107d-e; Walton 1997)

This new debate is part of an endeavour to improve the study of epistemology by investigating how the social interplay between agents promotes or inhibits the generation and transmission of knowledge. In the course of this lively debate, scholars like Goldman and David Coady became interested in applying such epistemological debates to concrete problem cases of societal concern, like how laypeople could find reasons to rationally trust some putative experts over others. More precisely, they focused on cognitive expertise, i.e. expertise about propositional knowledge. While this is a problem humans had to navigate for as long as there has been a division of labour, it gains in relevance in our modern, highly differentiated and complex societies. Especially in democratic contexts, people increasingly aspire to question their sources of knowledge critically. At the same time, general distrust in *experts* is rising in parts of society connected to populist political movements. On the other hand, from the 60s until the end of the century scholars in philosophy, psychology and sociology of sciences have focused on the limitations of science and expert judgments and the influence of monetary incentives on the objectivity of their work. Indeed, experts are neither infallible nor immune to corrupting incentives. Accordingly, the applied epistemological debate of the last two decades could be seen as an attempt to find a more nuanced approach to trust experts rationally. (see Chase and Coady 2019; Collins and Evans 2009, 143–45; A. Goldman and O'Connor 2021)

So far, however, scholars have mostly discussed cognitive expertise *in general* as if it were inconsequential for heuristics of what a domain specifically pertains to or how it is societally embedded. While many similar rationales are relevant in different kinds of expertise, not every heuristic is as effective in every domain. I will argue, instead, that heuristic strategies need to be adapted to different contexts to enable optimal orientation for laypeople. Having analysed the examples and case studies used in the philosophical debate and how heuristics have been

defended, I will judge that large parts of the debate have (at least) implicitly focused on scientific expertise, foremost in natural sciences like physics, medicine or climate sciences. Social Sciences or related non-scientific cognitive expertise have, except law, received scant attention. This is an unfortunate omission since such domains show distinct characteristics potentially relevant for the assessment of laypeople. Moreover, advice in such areas is very well demanded. Ergo, I want to analyse whether and how the general debate of the expert-identification question can be applied to the case of personal finance. The latter is of growing intimate concern to laypeople, noticeably distinct from usual case studies of the debate and attractive to malicious agents.

In this master's thesis, then, I want to analyse how laypeople could assess the trustworthiness of personal-finance advice. Accordingly, I will analyse heuristics which are likely to be affected by the shift to this specific domain and discuss whether they are applicable or how they need to be modified or concretized. Just like Goldman, I will assume the individual perspective of a layperson and discuss what she realistically could do to improve her epistemic position, excluding otherwise important political approaches. To not beg the question, I will assume that laypeople are constrained in money, time, energy, etc. which prevents them from learning the expertise themselves or engaging in overly complex methods to assess the trustworthiness of laypeople.

Additionally, I will use 'cheap talk' models of economic game theory to strategically analyse the relationship between experts and laypeople. To date, the philosophical expert-identification debate mostly adopted a passive perspective on this research question by discussing how laypeople best process the information they receive. We should, instead, also pay attention to how expert's reactions depend on the choices of laypeople, how agents are likely to anticipate behaviour and how this affects which actions are strategically sensible for laypeople. Moreover, this perspective could help in analysing how market contingencies affect the bargaining power of each participant. Fortunately, economic game theorists already addressed the strategic relationship of agents in asymmetric information transmission. Initially, Vincent Crawford and Joel Sobel were the first to propose cheap talk models where an expert can communicate information to a layperson but has a bias deviating from the layperson's interest. This literature and the debate following it, analyse the effects such biases can have on how much information experts will communicate, given that participants act rationally. I think it is promising to connect these two disciplines and complement the expert-identification question with heuristics that utilize strategic thinking. In doing so, I will spare the reader the mathematics involved in

these endeavours and translate the working of these modelling mechanisms into natural language. (see Ambrus and Lu 2014; Crawford and Sobel 1982; Farrell and Rabin 1996; Sobel 2013)

My thesis consists of two parts: In the first, I begin by clarifying preliminary questions relevant to framing and understanding the research question in Chapters 1 and 2. Subsequently, I will critically evaluate the expert-identification debate in Chapter 3, focusing on (strategic) heuristics of special interest in the second part. Consequently, In the second part, I will discuss what is peculiar about personal-finance expertise and analyse whether and how heuristics can be applied to this domain. More precisely, in the first chapter I will discuss what could be a sensible understanding of expertise. After all, to meaningfully talk about identifying proper experts, we need to agree on a common understanding thereof. In this vein, I will use Goldman's account as a convenient starting point to keep the debate connected to my research question. Philosophical accounts aside, also Harry Collins and Robert Evans' seminal sociological study of expertise will inform this endeavour. The Chapter itself is structured by four questions guiding the conceptual analysis, regarding whether expertise should be defined in a realist sense, relative to a social group, fulfil a sufficiency threshold and what exactly competency entails. The goal is to find a description that is close to the common understanding and relevant to laypeople in the context of the expert-identification question. (see Collins and Evans 2009; A. I. Goldman 2018; 2001)

In the second chapter, I will begin by defining what I understand under the term laypeople and what skills and resources we can expect them to have. This is relevant for keeping the heuristics I will propose realistic and appropriately targeted. Accordingly, I aim to address a broad public with limited means typical in 'Western' democracies. Subsequently, I will continue by discussing how fallible heuristics can rationally justify laypeople to trust putative experts. Especially C.A.J. Coady's pioneering work on the philosophy of testimony will guide this chapter. Starting with locating the contemporary debate in the larger research programme of social epistemology and philosophy of testimony, I continue with discussing the epistemological statuses of heuristics, risk and uncertainty. Finally, I will argue which kinds of information are accessible to laypeople and argue for the relevance of strategic thinking in this debate. (see C. A. J. Coady 1994)

In the third chapter, I will critically evaluate lay heuristics for the assessment of experts proposed by the general expert-identification debate. I will begin with heuristics that assess the incentive structure of experts and discuss their prevalence in the philosophical debate. Since

models of economic game theory also assess communication by strategically analysing each agent's incentives, this is the place where I introduce cheap talk models after Crawford and Sobel. I will finish this chapter by integrating the insights of such economic models into the general philosophical debate. The next heuristic I will discuss is Goldman's idea of assessing the past track record of experts using salient indicators of success where expertise post-hoc becomes 'exoteric', such as witnessing a successful rocket launch and inferring that those who worked on it must have had some special knowledge. (see Crawford and Sobel 1982; A. I. Goldman 2001, 106–8)

The third kind of heuristics I discuss in Chapter 3 tries to utilize the fact that multiple putative experts exist. Accordingly, I will discuss the epistemic value of consensus and agreement along with the conditions relevant for comparing experts' agreement on a testimony. After all, many scholars in the expert-identification debate share the intuition that propositions which are supported by many experts seem to be more robust and thus trustworthy, at least heuristically for laypeople. Thereby, I will connect the debate with Jacob Stegenga's insights on robustness-style arguments and the relevance of both independence and concordance of evidential sources. Finally, I will engage with game theoretic cheap talk models again, but this time with a modification that assumes one receiver interacting with multiple experts by Attila Ambrus and Shih En Lu. In this vein, I will examine whether laypeople can elicit a strategic advantage by comparing multiple experts even if both are biased. (see Ambrus and Lu 2014; Stegenga 2009)

With Chapter 4, the second applied part of my thesis will begin analogically to the first part with a characterisation of personal-finance expertise. Firstly, I will try to delineate the discipline and its practices by especially referring to Graber and Chatterjee's *De Gruyter Handbook of Personal Finance*. I will continue by describing the way laypeople usually encounter personal finance. In the third part of this chapter, I study how being informed by social sciences is affecting the testimony of personal-finance experts. Finally, I separately consider the practice of economic forecasting used by personal-finance experts and its epistemic value. (see Grable and Chatterjee 2022a)

Having a good understanding of personal-finance expertise, I assess how applicable the heuristics I discussed in Chapter 3 are to this specific disciplinary context. Starting again with incentive structures in the domain of personal finance, I will argue how the potential ambiguity and intrinsic uncertainty of this testimony affect the possibilities and heuristics of laypeople. Using Ben Almassi's conditions of detrimental incentives as guidance I further assess common incentive structures and institutional constraints for experts to bias their testimony. In the

process, I also investigate how game theoretically informed insights interact with the personal-finance context. Additionally, I discuss a modified cheap talk model by Roman Inderst and Marco Ottaviani that specifically analyses the practice of commission-taking regarding recommendations for financial products. I conclude this chapter by proposing another heuristic rationale inspired by Sanford Goldberg's paper: "If that were true I would have heard about it by now". (see Almassi 2017; Goldberg 2011; Inderst and Ottaviani 2012)

In the subsequent Chapter, I move on to the discussion of what I call Goldman's 'salient success' heuristics for assessing an expert's past track records and argue whether this strategy can be applicable to the context of personal-finance expertise. Eventually, I revisit the heuristics of comparing testimonies of multiple experts in the context of personal finance. Here, I discuss how consensus-building, and agreement are affected by the social science-informed peculiarities of personal finance and how independent or concordant testimonies can be here. In the end, I revisit strategic heuristics to induce truthful communication in scenarios of multiple experts again, following Ambrus and Lu's discussion of imperfectly informed experts, i.e. experts who receive noisy signals. Accordingly, I will ask how effective such heuristics can be but also investigate another way competition could benefit laypeople, introduced by Inderst and Ottaviani. (see Ambrus and Lu 2014; A. I. Goldman 2001, 106–8; Inderst and Ottaviani 2012, 405–7)

Other popular heuristics in the expert-identification debate would be, for example, heuristics of assessing the rhetoric and demeanour of experts, where some content-independent indicators could give hints as to whether an expert is performing well in a debate or engaging in manipulative behaviour. Even though associated rationales are deemed sizably fallible, they are still being entertained due to the general scarcity of indicators for trustworthiness. A more interesting, aspect I have mostly omitted in my thesis is the relevance of knowledge intermediaries and testimony chains in bridging the gap between experts and laypeople. This, for instance, relates to questions about how meta-experts certify experts for laypeople. Nevertheless, neither the rhetoric and demeanour of putative experts nor the rationales behind meta-experts and knowledge intermediaries significantly change the application to personal-finance expertise. Hence, I omitted them from my analysis.

In general, contributors to the expert-identification debate have not yet argued for an objective, complete list of possible lay heuristics, though the previously mentioned are the most frequently discussed. I think this can be explained by the fact that heuristics simply are strategies that give laypeople imperfect reasons to trust experts and thereby are not easy to

delineate or categorize. Hence, more heuristics could be imagined, and I do not claim to be exhaustive either.

Main part

1. A conceptual analysis of expertise

Before I can discuss strategies to identify personal-finance experts, we first need a common understanding of what experts in general are supposed to be. In common public discourse, people entertain a broad spectrum of associations with the term ‘expertise’. These conceptions appear to be very context-sensitive but seem to also share a family resemblance orbiting being ‘very competent’. Unsurprisingly though, different usages of ‘expertise’ conflict, which opens the potential for misunderstandings.

Therefore, I will start by proposing a general conceptualisation of expertise that intends to connect what most expert-associated phenomena share, thus staying connected to public discourse. Additionally, I want to focus on understandings of expertise that are relevant for laypeople wanting to improve their epistemic position. After having discussed the general understanding, I will address the specific case of personal-finance expertise in Chapter 4.

In conceptualising expertise, I will answer four questions and let them be the red thread of this chapter. The first is: Are experts endowed with a ‘real’¹ domain-specific competence or is it merely a social ascription or role²? Secondly, is expertise constituted relative to a contrast group like laypeople or expert peers? Thirdly, do experts need to fulfil a sufficiency threshold? Finally, what would this competency exactly entail?

A general point with broad agreement is that expertise always pertains to a field, speciality, domain, area, or discipline. This might be a thematic neighbourhood of actions the expert is competent about, or it is the function of a social role. Further, there need to be no fixed delineations between those domains, instead, there can be different ways of classification. For instance, expertise like physics, gravitational wave physics or astrobiology can overlap in different ways, which is a matter of convention informed by similarity of actions or

1 Also called objectivist (see Martini 2014b; Quast 2018; Scholz 2009) or realist (see Collins and Evans 2009; A. I. Goldman 2001; Grundmann 2024) in the debate.

2 Also called subjectivist (see Quast 2018), reputational (see A. I. Goldman 2001; Grundmann 2024), or relational (see Collins and Evans 2009) in the debate.

methodology or simply motivated by pragmatic reasons. Being an expert in all the aspects of physics without sacrificing depth, is, *ceteris paribus*, more involving than gravitational wave physics alone. (see Collins and Evans 2009; A. I. Goldman 2001)

I think it is useful to start with a discussion of Goldman's most recent account of expertise. Since he reinvigorated the expert-identification question in analytic philosophy and pioneered the discipline of social epistemology, his account of expertise received similar attention and became a central point of comparison for other participants in the debate. Note, that he focuses on cognitive expertise while leaving open how suitable his contributions could be for other types of expertise. Without providing a clear-cut distinction, he understands domains of propositional knowledge like sciences, humanities, or journalism as cognitive but not expertises like playing an instrument well or being a sports expert. (see A. I. Goldman 2001, 91; 2018, 4)

In the course of Goldman's engagement with the concept of expertise, he reformulated and revised his account several times without digressing too far from his main arguments.³ Here, I want to refer to his latest version, which he calls TL2⁴, involving two conditions:

[TL2] S is an expert about domain D if and only if (A) S has more true beliefs (or high credences) in propositions concerning D than most people do, and fewer false beliefs; and (B) the absolute number of true beliefs S has about propositions in D is very substantial. (A. I. Goldman 2018, 5)

While also being influential, these two conditions received specific criticisms. Before proceeding to them, I think we ought to discuss the general idea of conceptualising expertise as being constituted by real abilities in one's domain instead of, for instance, via a social role. (see A. I. Goldman 2001; 2018)

While Goldman and David Coady defend realism within a veritist (also 'truth-linked') framework, the condition of 'having true beliefs' is not the only way of establishing a realist account of expertise. Collins and Evans in their sociological work are also starting from the premise that there are 'real' experts "(...) who know what they are talking about" (2009, 2), and classify different kinds and levels of expertise. They do not reduce expertise to some necessary or sufficient conditions but distinguish levels of skill and the social environments

3 In his earlier 2001 paper, Goldman also included as a third (C) condition the need to have the ability to exploit one's true beliefs to find new true beliefs to new questions of one's domain, which he later abandoned. (see A. I. Goldman 2001, 89–92; 2018, 3–6)

4 TL stands for 'truth-linked', while '2' denotes that this is the second, and last, modification made within his paper. (see A. I. Goldman 2018)

needed to acquire them. Watson calls this a *process* account in contrast to ‘fixed-criteria’ accounts like those of Goldman and David Coady, who employ classical conceptual analysis with necessary or sufficient conditions. (see D. Coady 2012, 28–30; Collins and Evans 2009, 13–44; A. I. Goldman 2001, 89–92; 2018, 3–6; Watson 2020, 5)

Collins and Evans use the term expertise also for ubiquitously possessed skills if they require a sufficient amount of learning, hence calling them *ubiquitous* expertises. This, they see as a demand of a realist approach, where competence but not its social prevalence ought to be important. To make their point, they use the example of English and French speakers in France, where French should not be seen as less of an expertise, even though most people in this local region possess it. (see Collins and Evans 2009, 13–23)

While this is a sensible intuition, I think that accounts which define expertise as special competence in comparison to a general lay audience can be realist just as well. As Goldman or David Coady both did, one can include a comparative condition such as ‘being more competent than most people’ and still require that this competence is based on real qualifications in a domain. The sensibility of this approach hinges on the usefulness of reserving the term expertise for skills that are held only by a small part of the population deriving from an epistemic division of labour where people depend on each other’s specialisation. Collins and Evans, instead, use the term ‘specialist expertise’ to contrast knowledge in esoteric domains from ubiquitous ones. What Collins and Evans call specialist expertise, is equivalent to what Goldman or David Coady reserve the term ‘expert’ for. (see Collins and Evans 2009, 13–44; A. I. Goldman 2018, 3–7; D. Coady 2012, 28–30)

Luckily, Collins and Evans’ clear classification allows comparisons with other accounts without preliminarily deciding the question of whether expertise-ascriptions should depend on their prevalence in society. Alternatively, non-realist accounts of expertise might be motivated by difficulties in clarifying the performance and scope of different expertises, such as Agnew, Ford and Hayes’ evolutionary selection theory of expertise or Turner’s public influence account. Another author often cited in this context is Matt Stichter, who emphasizes the difference between ‘having expertise’ and ‘ascribing expertise’ to others. The latter is always a social decision that does not need to succeed in targeting the true experts. I am very much aware of this distinction as in my paper I am discussing the best but still fallible ways to ascribe expertise from a lay perspective. For that, however, I first need a theory of what a true expert would be. Goldman disagrees that we should use the term ‘expert’ reputationally for those who have imperfectly been ascribed expertise. I also think that it will be less prone to

misunderstandings when we call those who we think are real experts ‘experts’ and add an adjective like ‘supposed’, ‘so-called’ or ‘putative’ to the noun ‘expert’ when we want to highlight the mere (unjustified) ascription onto others. (see Collins and Evans 2009, 13–44; A. I. Goldman 2018, 3–7; Stichter 2015; Watson 2020, 141–53)

I want to argue that realist views are the most useful for thinking about the expert-identification question. Firstly, such a position avoids the problem of constructivist and reputational accounts since real expertise neither requires public awareness nor approval. Secondly, it would give laypeople clear, normative reasons to seek expert advice. Goldman emphasises that even if no expert could fulfil this wish, only the laypeople’s belief that experts can improve their epistemic situation could motivate their advice-seeking behaviour. I think this argument has a strong intuitive appeal as there would be no rational reason for a layperson to find people with the reputation of being an expert apart from this reputation making it more likely that they are truly experts. Only in the latter case, their testimony can improve the epistemic position of the layperson. Moreover, there are uncontroversial examples of people exhibiting real expert skills like car mechanics, surgeons, or taxidermists. This might not be as clear in every supposed field of expertise, but simply requires a case-by-case analysis instead of a sweeping rejection. (see Collins 1992; A. I. Goldman 2018, 6–7; Scholz 2009, 189–91; Watson 2020, 49–53)

We can understand the fallibility and interests of financial advisers, lawyers, politicians, art and literacy critics, doctors, builders, car mechanics, and travel agents without concluding that they are not more expert in their area than we ourselves. Neither anarchy nor nihilism follows from the recognition that there is no magical escape from the pangs of uncertainty that underlie our decisions. (Collins 1992, 167)

Following these considerations, I think we should reserve the term ‘expert’ for those who have some real competency in their domain. Some of these disciplines, however, might also fulfil a social service function. As Goldman understands ‘expertise’ as a multi-faceted phenomenon, he is sympathetic to the existence of multiple sensible definitions of this concept. In this vein, he himself proposed the following functionalist notion: (see A. I. Goldman 2018, 3–5)

[CAP] S is an expert in domain D if and only if S has the capacity to help others (especially laypersons) solve a variety of problems in D or execute an assortment of tasks in D which the latter would not be able to solve or execute on their own. S can provide such help by imparting to the layperson (or other client) his/her distinctive knowledge or skills.⁵ (A. I. Goldman 2018, 4)

First of all, this approach is compatible with Goldman's veritist TL2 account regarding the underlying realism but instead of characterising expertise by true beliefs, it does so by the role she needs to fulfil. It is also susceptible to the same argument on whether a comparative condition can be realist. Christian Quast followed a similar approach in describing expertise by its function in society and the necessity of having real expertise, as agents being sufficiently disposed to fulfil their function. He argues for the intuitiveness of this functionalism by referring to common language use, depicting experts as needing to be more competent than laypeople in contrast to Collins and Evans. Quast additionally argues that any kind of expertise ought to be characterised in this way and that this concept not only involves the ability to help laypeople but also the willingness to do so honestly. In this sense, an agent with real skills should not be called an expert but just competent. (see Collins and Evans 2009, 13–23; A. I. Goldman 2018, 3–5; Quast 2018)

Against a service-orientated notion of expertise, one could argue that not every expert interacts with laypeople. We could, for example, imagine a mathematician who is not interested in public issues and solely engages with her abstract study and her colleagues. Would it not be strange to deny this person's expertise? Michel Croce further adds that not every expert has the pedagogical or communicatory skills to help laypeople. Vice versa, he continues, some agents might be able to teach something in a domain without even being a top performer. (see Collins and Evans 2009, 2–9; Croce 2019, 12–19; Grundmann 2024, 1–5; Watson 2020)

I would like to respond, that we could still see the mathematician as serving the public indirectly by her insights that might benefit the discipline and in turn support other useful developments. Alternatively, this expert could be seen as interacting with the lay public by a chain of testifiers. It is not uncommon, that some experts collect a discipline's generated knowledge and teach it to intermediaries like science communicators or educators who in turn render it comprehensible for laypeople. While this would also explain why even detached experts are getting paid by society, one might construct an extremer example of a reclusive expert who, after having learned the expertise, exerts it far away from society. This eremitic

⁵ CAP refers to capability, which means that experts are defined by their capability to help others regarding matters of their discipline. (see A. I. Goldman 2018, 3–5)

expert, though rather odd, may indeed be an example where an agent intuitively retains her expert status. (see Collins and Evans 2009, 2–9; Croce 2019, 12–19; Grundmann 2024, 1–5; Watson 2020)

Therefore, I am not going as far as Quast in deeming it necessary for all experts to be service-oriented. Moreover, I think we should not include ‘being trustworthy’ in what makes one an expert. As with everyone else, experts can be moved by many motives. Some of these can occasionally incentivise misconduct but would not destroy their ability to help forever. In Quast’s deontic interpretation, an agent could repeatedly become and cease to be an expert depending on her dishonest inclinations but independently of her competence. I think it is rather common to both speak about diligent and untrustworthy experts. Instead, I have introduced the notion of service-orientation here not because I think that any expert must have it but very well every kind that is relevant for laypeople in the context of the expert-identification question. A tax consultant, secluded from society, could not employ her expertise without serving clients. Conversely, no layperson would need to decide on trusting an expert she cannot learn from anyway. This follows for scientific expertise just as well as for services like medicine or personal finance. Yet especially consultancy expertises are necessarily service-oriented. (see Quast 2018)

Accordingly, what laypeople should try to identify are experts as epistemic authorities. Most famously, this notion has been described by Linda Zagzebski. In her work, she builds upon the idea that testimonies can be a legitimate source of knowledge and rational inquiry is not in conflict with critically trusting others as I will describe in Chapter 2.2. Without describing her work in detail, she specifically analyses the role of expert testimony and whether non-experts are justified in accepting those. If someone is an epistemic authority, and the non-expert recognises this, then the latter is not only justified but even rationally required to believe her testimony. Croce further refines and connects this idea to the expert-identification question by comparing it to Goldman’s functional notion: (see Croce 2019, 12–16; A. I. Goldman 2018, 3–5; Martini 2020, 115–16; Zagzebski 2012)

On the service conception of epistemic authority, I adopt, following Raz (1986, 2006), Zagzebski (2012), and Jäger (2016), a subject is an epistemic authority for another insofar as they can help their interlocutors achieve some epistemic goal in a given domain through their superior knowledge and/or understanding to the interlocutors. Thus, the function of epistemic authorities on the service conception is just a slightly revised version of Goldman’s NO-CAP.⁶ (Croce 2019, 14)

⁶ Croce refers here to Goldman’s functional definition of expertise I cited earlier but calls it NO-CAP to further differentiate that the aim of the function are novices, thus NO. (see Croce 2019, 3 Footnote 3)

For my research question, it is not important to immerse ourselves in the ramifications of this concept but just to see experts as agents who laypeople can, *ceteris paribus*⁷, rationally trust within the scope of their field, hence from whom they can learn or otherwise epistemically profit. (see Croce 2019, 12–16; Zagzebski 2012)

So far, I argued that for the expert-identification question the laypeople's goal is to find experts who are competent in a domain and can help them in this regard. This entails that experts need to be more competent than laypeople or else they could not support them. The required level of expertise, then, depends on the skill prevalent in these two groups, relative to time and location.⁸ Should such experts, however, also fulfil a non-comparative, minimal threshold of competency? Goldman argued that requiring such a threshold would prevent us from also ascribing expertise when 'experts' know little about their field just because the rest of society knows nothing about it. Goldman calls this the 'new domain' intuition as the common example for such scenarios are infant sciences, engaging in early, very uncertain efforts. David Coady and Croce do not share this intuition and question why we could not simply see someone as an expert independent of the absolute level of competence. They argue that as long as experts are significantly more competent than laypeople, they can inform, teach or otherwise epistemically support the latter and thus fulfil the conditions of being experts as epistemic authorities. I, too, think the latter approach better describes what laypeople care for in the expert-identification question. Accordingly, I see experts and laypeople on a dynamic range of competence, where a significant epistemic asymmetry is enough to separate them. (see D. Coady 2012, 28–30; Croce 2019, 5–7; A. I. Goldman 2018, 3–6)

Finally, I want to return to the question of what the *competence* of experts more precisely entails. Goldman's truth-linked account requires experts to have more true beliefs and less false beliefs than most people. In response, David Coady argues that people plausibly can have less false beliefs, when they are simply disinterested in the question compared to an expert who engages in fallible discourse and argues that simply the condition of having more true beliefs is preferable. Oliver Scholz further adds that laypeople can also have more true beliefs than

7 One reason to insert a *ceteris-paribus* clause, is that contrary to Quast, experts as epistemic authorities can have reasons to be dishonest, in which case a layperson might justifiably refrain from trusting here, despite her proper ability to help. (see Quast 2018)

8 Alternatively, one could see the state of the art of an expert community as a benchmark for expertise, but this would not conform with my conceptualisation of experts as epistemic authorities. (see Collins and Evans 2009; Watson 2020)

experts when they extensively read up on the more established facts of the domain., while the experts engage in messy frontier work. Goldman, however, specifically understands his condition as having a better ‘ratio’ between true and false beliefs, which is often overlooked by replicants in the debate. Notwithstanding this ratio, it remains questionable whether true beliefs best describe expert competence. Collins and Evans, for example, highlight the importance of tacit knowledge, which experts learn by social enculturation in their expert community. This, for example, includes relevance and plausibility judgments entrenched in the expertise. (see D. Coady 2012, 28–20; Collins and Evans 2009, 2–27; A. I. Goldman 2018, 3–6 esp. footnote 3; Scholz 2009, 192–94; Watson 2020, 53–57)

Other scholars proposed alternative epistemic desiderata instead. Goldman himself is open to replacing true beliefs with justification or evidence but warned of thereby introducing relativistic notions. While Goldman considered such options only superficially, Croce, Scholz and Thomas Grundmann more extensively entertained such concepts. Croce and Scholz, for instance, argued that true beliefs could be the result of mere memorisation, whereas experts are more than knowledge depositories. Instead, experts have an ‘understanding’ of their field, which involves knowing the systematic connections between propositions and their explanatory relevance. (A. I. Goldman 2018, 5–6; Grundmann 2024, 8–10; Scholz 2018; 2009, 192–94; Watson 2020, 82–85)

Be it evidence, understanding or ‘justified, true belief’, this aspect of conceptualising expertise is derivative of the old epistemological question of what it means to know something. In this thesis, I cannot meaningfully add to this debate and thus I will stay agnostic about the exact conceptualisation of knowledge. To identify experts as epistemic authorities by the appropriate use of heuristics it is inconsequential whether competence entails true beliefs or understanding. Hence, in the rest of the paper, I will simply refer to the expert’s ‘knowledge’ or ‘competence’.

In conclusion, when I speak about experts in this thesis, I mean agents that are epistemic authorities for laypeople since they are significantly more competent than laypeople within a domain, given the respective time and region. This also involves a social-service function that such experts would fulfil. Moreover, no minimal absolute threshold of competence is necessary since all that matters is that laypeople can rationally defer to these agents. Real experts, by definition, are adequately competent in a realist sense but do not need to be trustworthy since they can be affected by corrupting incentives.

2. Laypeople and their epistemic justification to trust experts

2.1 Who are the laypeople?

After having argued that, for this paper, I am interested in experts as epistemic authorities for laypeople⁹, I now want to characterize the opposite side of this relationship. By implication, laypeople are those who do not fulfil the conditions of being an expert, which are most of us in most contexts. Nevertheless, I think it is sensible to give some more characterisation as it clarifies which kind of heuristics are suitable for them.

Lay people are constrained in time, energy, and previous education. It would be cynical to answer the expert-identification problem by demanding that laypeople should just learn and become experts themselves. I would even go further and claim that we should not presuppose anything more than a regular ‘high school education’. Or rather, the kind of general school education most people in richer countries have, including the ability to properly write, read, calculate, and what Collins and Evans call ‘ubiquitous expertises’, i.e. basic abilities people learn through social enculturation in their society and are necessary to navigate it. This includes things like natural language speaking, customs, and even the ability to superficially distinguish political parties and their values. (Collins and Evans 2009, 13–15)

Still, it might be questioned why we would not ask laypeople to do a moderate amount of research into the content of a discipline to gain some critical distinction. Firstly, lower to moderate investments into a discipline’s content are hardly enough to answer one’s questions and neither help with distinguishing whom to trust. After all, one still lacks a lot of understanding, so one cannot be certain not to miss something that would change the way one judges who an expert is in the field. Secondly, superficial information could instil a wrong sense of certainty, as people tend to be systematically over-confident regarding uncertainty in general but also concerning the incompleteness of their knowledge. Thirdly, doing so would require a lot of time and effort. Since I want to find solutions for a broader societal problem, I should not presuppose time most people who have to work and care for others do not have. Moreover, people are confronted with a multitude of questions every day where expert knowledge would be useful. One cannot become an expert in everything. Therefore, the heuristics I am discussing

⁹ Novice, non-expert and layperson are used synonymously in the debate following Goldman, except for Scholz who proposes to use ‘novice’ for advanced laypeople that are in the process of becoming an expert. (see A. I. Goldman 2001; Scholz 2009, 191)

in this thesis need to be sensible to these time constraints and I cannot presuppose any special knowledge beyond the basic competencies people learn in their national schooling system. (see Angner 2006, 6–7; Collins and Evans 2009, 13–15; Morgan and Mellon 2011)

Is it, however, not often the case that laypeople use specific bits of knowledge to check the plausibility of expert testimony? I think it is very well imaginable that laypeople can use fragmentary specialised knowledge to assess expert testimony, but such considerations are very heterogeneous and contextual. Moreover, a good level of general education could instil a better understanding of institutions and their relationships of interests, calibrating the layperson towards trusting better knowledge intermediaries, which could raise the chance of productively using expert testimony. Although interesting, discussing the use of partial specialist knowledge would require its own dedicated paper and would be a digression from the more usual case of utilizing expert testimony without any specialist insights. (see Lane 2014, 100–105)

This is not to say that all laypeople share identical resource constraints. Some might have more time available, enjoy a more extensive general education or possess more money. These might allow them to follow heuristics more extensively or use more demanding ones. Conversely, rather disenfranchised laypeople might have even less time to engage with expert testimony, because, for example, they might need to work longer hours, more strenuous jobs, have more care responsibilities or cannot pay potential advising fees. As a result, they might have to either blindly trust or forgo expert advice completely. Here, ethical concerns enter the epistemic debate of expert testimony. This opens up potential for debates on justice and could connect the expert-identification question with debates on ‘epistemic injustice’, popularised by Miranda Fricker. (see Grasswick 2018, 4.1; Fricker 2007)

Depending on the social group affiliations of both experts and laypeople, biases and aversions could, further, affect the testimony transmission. For example, Boudewijn de Bruin discusses gender and racial disparity in financial markets and argues that members of disadvantaged groups often receive inferior financial services. He addresses, for instance, that women pay higher interest rates for mortgages with worse conditions, that their loans are more frequently rejected, and that they face less favourable treatment in case of defaulting than men. Both, thinking about specific heuristics for different resource backgrounds and specific heuristics addressing particular social identity biases would distract from the more general, and flexibly applicable, arguments in the expert-identification debate. Instead, I hope my previous characterisation of laypeople is useful in most cases. (see de Bruin 2021; Grasswick 2018, chap. 4.1)

In framing my endeavour like this, I am in good company as these presuppositions have primarily guided the debate since Goldman.¹⁰ This further entails that laypeople cannot directly assess putative expert testimony without using heuristic judgments. In the following section, I will discuss how laypeople can be justified in trusting a putative expert without using domain-specific knowledge. (see A. I. Goldman 2001, 89–92; Grundmann 2024, 11; Goodwin 2011, 288–91; Martini 2020)

2.2 Rational trust in epistemic authorities by employing fallible heuristics

The expert-identification question is an applied debate on the epistemology of testimony. A central question in the latter is, whether the testimonies of others can be a justifiable and fundamental source of knowledge (anti-reductionism) or whether believing in testimonies can only be indirectly justified by other epistemic faculties of the testifier like natural sense perception (reductionism). (see C. A. J. Coady 1994)

In the 19th century individualistic enlightenment philosophy, scholars like John Locke were sceptical towards testimonial sources.¹¹ They saw reference to testimony in conflict with the autonomous use of one's rational faculties, akin to uncritically parroting rumours. The scrutiny towards absolute authorities inspired them to view social structures as rather biasing and corrupting than helping knowledge generation. Contrary to focusing solely on the oppressive potential of human interaction, C.A.J Coady questions the principled distinction between testimonial and other sources of knowledge. David Hume, for instance, saw testimony as relayed inductive inference since testimony can be discovered as unreliable by observation. C.A.J. Coady responds that testimony can also be questioned by other testimony without personal observation, for instance, via credentials or cross-examination. Furthermore, testimony can move us to reject personal observation, for example, when we know we are hallucinating. (see C. A. J. Coady 1994, 79–100; Hume 2007, sec. 10; Locke 1979)

10 Notable exceptions are, for example, Saúl Pérez-González & María Jiménez-Buedo and Lane who advocate for differentiating more between different kinds of laypeople based on their competencies or abilities. (see Anderson 2011, 144–47; Lane 2014, 111–13; Pérez-González and Jiménez-Buedo 2023)

11 A quote from Locke represents this individualistic attitude well and is cited often in the contemporary debate: “The floating of other men's opinions in our brains makes us not one jot the more knowing, though they happen to be true. What in them was science is in us but opinion; and yet, whilst we give up our assent only to reverend names, and do not, as they did, employ our own reason to understand those truths which gave them reputation. Aristotle was certainly a knowing man; but nobody ever thought him so because he blindly embraced and confidently vented the opinions of another. And if the taking up of another's principles without examining them made not him a philosopher, I suppose it will hardly make anyone else so.” (Locke 1979, bk. 1, chap. 3, sec. 24)

Moreover, without treating testimony as potentially fundamental, we are not able to properly deem information from memory or about history as knowledge, after all, memory is a recreation of past content from one's brain and historical knowledge is impossible without testimonial transmissions. Even more so, any bigger human project requires trust and dependence on colleagues, making it difficult to explain scientific knowledge without accepting testimony as a possible fundamental source. Such 'first principles' are incapable of direct proof but can produce fallible beliefs, just as sense perceptions can. A difference, C.A.J. Coady admits, is that the reliability of testimony can be intentionally decreased by the testifier. (see C. A. J. Coady 1994, 124–42)

This debate also marks a shift from an individual perspective towards social epistemology, which opens the focus to the social cooperation necessarily involved in knowledge generation and transmission. In a complex society with high divisions of labour also the epistemic work is divided among many different experts. To orient oneself and make smart decisions one often requires expert advice in particular domains.¹² (see A. Goldman and O'Connor 2021)

So, if one agrees that experts as epistemic authorities exist, that testimony can be a source of knowledge, that due to lack of time and competence laypeople cannot directly assess expert testimony, and yet such domain-specific knowledges would be very relevant for them, it follows that laypeople are epistemically dependent on experts. Prominently, the epistemic dependence of laypeople has been contended by John Hardwig and described as the fact that laypeople need to blindly trust experts if they want to use their knowledge. After all, how could a non-expert differentiate between expert testimonies while lacking access to the epistemic processes that constitute such specialised knowledge?¹³ (see A. I. Goldman 2001, 86–89; Hardwig 1985; Martini 2020, 115–16)

12 Interestingly, similar concerns also arise within expert communities. In complex areas of (scientific) expertise, working groups often contain participants with different specialities whose knowledge in neighbouring fields fades quickly with distance. Since the cultivation of specialised expertise requires a lot of time, effort, and personal teaching by peers, agents cannot become experts in everything. Hence, they often need to cooperate with neighbouring disciplines to make progress. Moreover, the scope of projects alone often requires the participation of many individuals. David Coady and Lackey even argue that cooperation is not just often a necessity but also carries epistemic virtues, like enabling them to see evidence from new perspectives. Monitoring by peers might, further, incentivise experts to be more conscientious. (see D. Coady 2006, 68–72; Collins and Evans 2009, 91–112; Collins 1992; Lackey 2013)

13 In this vein, Walton cites Sextus Empiricus' doubt in resolving this issue: "Who is to be the judge of skill? Presumably, either the expert or the nonexpert. But it cannot be the nonexpert, for he does not know what constitutes skill (otherwise he would be an expert). Nor can it be the expert, because that would make him a party to the dispute, and hence untrustworthy to be a judge in his own case. Therefore, nobody can be the judge of skills." (cited after Walton 1997, xiii; see also Bekker 1842)

Without wanting to settle the debate on reductionism in testimonial knowledge, Goldman emphasised that, to the contrary, there are certain kinds of reasons that justify non-experts well enough to increase or decrease trust in (expert) agents. Goldman presents the following thought experiment:

As you relax behind the wheel of your parked car, with your eyes closed, you hear someone nearby describing the make and color of the passing cars. Plausibly, you have prima facie justification in accepting those descriptions as true, whether this prima facie entitlement has an a priori or inductivist basis. But if you then open your eyes and discover that the speaker is himself blindfolded and not even looking in the direction of the passing traffic, this prima facie justification is certainly defeated. So what you empirically determine about a speaker can make a massive difference to your overall justifiedness in accepting his utterances. (A. I. Goldman 2001, 88)

This example ought to give intuitive credit to the idea that the trustworthiness of testimony can be assessed indirectly. Here, not being able to observe other cars ought to be analogous to a layperson's inability to assess expert evidence and discovering that the other person is not looking at the cars refers to observing that the other is no expert herself, ergo being not trustworthy.¹⁴ (see A. I. Goldman 2001, 87–89)

Instead of assessing expert testimonies directly, laypeople can use indirect heuristics to assess the trustworthiness of expert testimony. These are simple rules of thumb that ought to make a good judgment more likely. In my thesis, I follow Gigerenzer's *fast and frugal heuristics* program that understands heuristics as rules which are epistemically beneficial as they are adaptive to constrained epistemic facilities, environments, and time. This contrasts with the understanding of Daniel Kahneman who investigated biases and heuristics as habitual but less than rational deviations from proper inferences. Since I am engaged in a normative research question, my aim is not to analyse which heuristics are used in trusting experts but rather which of those would be epistemically optimal, given laypeople's constraints discussed earlier. A further parallel to the fast and frugal heuristics program is their context dependence. Each heuristic has epistemic advantages and pitfalls that need to be evaluated given the situation. Their justification might differ between lay person, field of expertise, available

¹⁴ Certainly, this example is simplified as observing the colour of cars is no esoteric expert skill and thus observing that the testifier is not looking at the traffic is a rather easy heuristic to learn that the testifier had no access to the source of knowledge, entailing that the driver has no justification to believe the testifier. (see A. Goldman and O'Connor 2021, 87–88)

resources, etc. Nevertheless, I am going to argue that some central rationales appear in most heuristics. Jutta Schickore further summarizes: (see Kahneman 2012; Wheeler 2020, chap. 7)

“Heuristics-based methodologies of discovery are neither purely subjective and intuitive nor algorithmic or formalizable; the point is that reasons can be given for pursuing one or the other problem-solving strategy. These rules are open and do not guarantee a solution to a problem when applied.”(Schickore 2022, chap. 8.3)

These characterisations resemble ‘inferences to the best explanation’(IBE), which Goldman and Boaz Miller connect to the expert-identification-debate. This non-formal, heuristic inference treats explanatory merits like parsimony or scope as likelihood indicators. Depending on an agent’s epistemic environment, empirical observations can make certain explanations more likely. Miller uses the example of fingerprints in a crime scene that do not necessarily prove an agent’s involvement but make it, *ceteris paribus*, more likely within an incomplete set of information. All the clues weighed up might point in one direction and justify a judgment as *best* explanation. In a similar vein, Melissa Lane emphasised that judgments do not entail knowledge but are decisions to believe given incomplete information. (see Miller 2013, 1299–1300; Lane 2014, 98–99; Schickore 2022, chap. 8.3; Wheeler 2020, chap. 7)

If heuristics are not formal methods of proof, what criteria do they adhere to? I would argue that a necessary condition for a *useful* heuristic is to lead to good decisions more likely than chance, more often compared to not using an intentional decision-making device or to elicit a higher quality in judgments than without. The usefulness of a heuristic, then, depends on its accuracy plus the time and effort it demands. Similarly, one might distinguish between the prospect of gaining new justified beliefs by testimony and avoiding error. Here, one can think about a matrix of true and false positives along with true and false negatives and think about epistemic attitudes that answer how the risk of trusting the wrong testimonies is to be balanced with the risk of not trusting the right testimony. C.A.J. Coady calls the preference for avoiding error ‘epistemic caution’ and for maximizing truths ‘epistemic boldness’. Which attitude is more sensible depends on the contextual consequences for the affected agents. After all, the amount or strength of reasons for convicting someone for a crime ought to be more demanding than the reasons we need to be justified in buying a good washing machine. Since human lives involve many projects that compete in energy and time, I want to remind us that it is not feasible to treat every decision as equally serious. Sometimes we need to relax our epistemic standards, which is unproblematic if we choose the right ones. (see C. A. J. Coady 1994, 109–13)

So far, I have used the term *risk* in its colloquial sense, but especially concerning assessments of expert testimony, *uncertainty* is the more frequent problem. Especially in the context of Bayesian updating, uncertainty refers to the likelihood of an event where, *prima facie*, no information on frequencies is available, whereas risk more commonly refers to random events where they are known, like in classical dice throws. Uncertainty is inevitably involved in the use of heuristics due to the incompleteness of empirical evidence and internal consistency. Also, expert testimony itself often involves uncertainty, which is consequential for lay heuristics that assess them. Lane discusses three varieties of those: Firstly, uncertainties intrinsic to the phenomena being studied, i.e. frequency fluctuations in weather systems that no method of study can predict. Secondly, uncertainties due to imperfections in the present methods of science, i.e. modelling choices, the selection of parameter and confidence intervals, and constraints in data collection. Thirdly, competitive uncertainties arise from disagreements between different experts, disciplines, and research cultures. (see Lane 2014, 105–11)

I think Lane’s distinction between intrinsic uncertainties and those conditional on present methods of science is not definite. Scientist’s knowledge of whether an uncertainty is truly intrinsic to the nature of a phenomenon is mediated by methods of science. It is imaginable that events that seemed uncertain in the past can be well predicted and described with newly developed methods. Perhaps Lane meant to distinguish between radical uncertainties¹⁵ and those that could be reduced by research, which nevertheless is difficult to predict. From the perspective of a layperson, however, these two kinds are indistinguishable anyway. Without having insights into how scientific methods work and empirical data presents itself, a layperson could not know whether uncertainty is inevitable or just a matter of temporary scientific limitations. Since laypeople’s reason to defer to expert authority is not to generate knowledge themselves but to decide on the best available evidence, the difference between inevitable and methodological uncertainty is inconsequential. Taking the epistemic perspective of laypeople, I will from now on refer to Lane’s first two varieties both as ‘intrinsic uncertainty’. (see Lane 2014, 105–11)

Unfortunately, laypeople cannot even distinguish this intrinsic uncertainty from the third variety, competitive uncertainty. Lane writes that the latter disagreements often arise from the first two varieties, i.e. ‘intrinsic uncertainty’. I would even argue that disagreements which do not arise from these intrinsic uncertainties can only be the consequence of epistemically

¹⁵ This term has been described by Mervyn King and John Kay as uncertainty which cannot be resolved by any kind of inquiry. (see King and Kay 2020)

unjustified social reasons, since only those concern the ‘content’ of the expertise. A group of scientists might, for example, defend an outdated position they had developed even though evidence decisively disfavours it by now to save their reputation of being pioneers. However, from the perspective of a layperson, every disagreement between experts looks alike. Without technical insight, they, *prima facie*, are unable to distinguish ‘normal’ uncertainty from epistemically detrimental disagreement. The expected intrinsic uncertainty, moreover, is an important factor in whether some heuristics work well, which generates an additional epistemic burden on laypeople. Judgments in the presence of uncertainties are, however, challenging as both experts and laypeople are prone to systematic overconfidence, empirical evidence shows. (see Angner 2006; Lane 2014, 105–11; Morgan and Mellon 2011)

Uncertainty aside, which kinds of heuristic evidence could be accessible to laypeople? Following the expert-identification debate, they should be interested in two aspects: Firstly, is the putative expert competent regarding my question? Secondly, is the expert inclined to honestly testify to her best knowledge? Any kind of heuristic needs to address either or both questions. According to Collins and Evans, laypeople typically possess the ability of ‘ubiquitous’ discrimination as part of an ‘ubiquitous expertise’, which they learn due to their social embeddedness in a society, which allows them to (epistemically) orient themselves in it. Ubiquitous discrimination, further, is a lay variant of a ‘meta-expertise’ consisting of social judgments about whom rather than what to trust. Harry Collins and Martin Weinel continue the endeavour to describe how laypeople use expert testimony by calling this process ‘transmuted meta-expertise’. Using social facts that are accessible to laypeople, they can transmute social judgments into technical judgments to decide on technical questions without needing the expertise to do so directly. (see Collins and Evans 2009, 13–23, 45–52; see Collins and Weinel 2011, 401–7)

Such social facts are, for example, the location of a putative expert in the social network of experts. This could, for instance, mean learning whether the testifier has a PhD or university position in a domain and what that means in relation to other experts. Also, other credentials, the presence of corrupting incentives or the appropriate demeanour of a putative expert are social facts open to laypeople. These socially transmuted judgments, though still fallible, give laypeople reasons to trust experts that are better than chance. (see Collins and Weinel 2011, 401–4; A. I. Goldman 2001, 87–89; Goodwin 2011, 288–91)

How effective it is to transmute social knowledge depends on how closely connected the layperson is to the social world of the respective experts. The better socially embedded she is,

the more she can know about social rules of expertise accreditation and what they mean. This kind of understanding is often tacit, in that their possession is often not salient and difficult to communicate. In its intense form, intimate social knowledge enables local discrimination regarding the trustworthiness of putative experts, like remembering self-experienced breaches of trust or demonstrations of virtue. Nevertheless, such transmuted judgments are prone to bias and not always reliable. (see Collins and Evans 2009, 48–52; Collins and Weinel 2011, 401–7)

Finally, next to social facts, I want to consider another type of information accessible to laypeople, which I call ‘basic information’. These are simple facts one can easily find or verify via a quick internet search. For example, how expensive a putative expert is, how often a paper has been cited or which social group an expert belongs to. Importantly, I do not argue that laypeople are likely to know what that means, just that such information is very unambiguous. Laypeople can use public, non-complex information since if they were wrong others would likely already have corrected them. The easy verifiability of such basic information makes them trustworthy.¹⁶ The interpretation of basic information is supported by our general education. Laypeople can, moreover, easily observe which expert community claims to be able to answer which set of questions. For instance, this is like asking a car mechanic to fix your car, asking a physicist about the basic laws of nature or an economist about the development of the interest rate in the coming year. However, it is much more difficult to assess whether experts should be trusted in this regard and by how much. For instance, later in this thesis, I will cast doubt on the ability of personal-finance experts to predict economic developments in the more distant future. (see Anderson 2011; Sobel 2013, 307–8, 327–28)

Contrary to most contributions to the philosophical debate so far, I will argue that such heuristics do not need to be constrained to passively analysing hints for trustworthiness but can also strategically induce actions moving experts to be honest. By using game theoretic tools, I will analyse how the social positioning of each party, and type of expertise can change the playing field of assessing expert testimony. Hence, I will show that the expert-identification question is treated incompletely without considering the market dynamics that are involved in the expert-lay relationship. (see Goodwin 2011, 293–95)

¹⁶ Without going into detail, the trustworthiness of ubiquitously available and easily verifiable information is also supported by game theoretic ‘disclosure games’ that are another type of communication model next to cheap talk games, which I will first address in Chapter 3.1.2. Paul Milgrom analysed models of strategic communication between laypeople and experts and concluded that the mere possibility of verifying a proposition incentivises honesty by experts even though they would otherwise prefer not to be honest. If laypeople can find out anyway there is no reason to fool them, though strategically withholding information can still be sensible for experts. (see P. R. Milgrom 1981; P. Milgrom and Roberts 1986; Sobel 2013, 307–8, 327–28)

An example of a heuristic judgment that utilises strategic thinking would be the following: Laypeople could question how an expert's testimony would affect the strategic relationship of experts, laypeople, and (market) participants if it were true, and other people (eventually) learn this testimony. If the anticipated consequence does not seem to be the case, then this might give one a reason to doubt the testimony. Relatedly, Goldberg wrote a paper on the rationale: 'If that were true, I would have heard about it by now'. In short, this heuristic judgment refers to the likelihood of the testimony reaching the agent from other sources if it were true. Following Goldberg, this depends on the newsworthiness of the testimony and whether the agent can expect reliable coverage within her epistemic environment. More precisely, the argument hinges on how well an agent is epistemically positioned, i.e. how likely and fast information reaches her, and how important the news is being deemed by one's epistemic peers and sources for relaying the information. If the information were true and newsworthy and the agent is epistemically-reliably covered then it is likely that the testimony would have already been received faster or from specific other sources or multiple different sources.¹⁷ Moreover, such newsworthy testimony might be perceived in public discourse and be partially common knowledge. (see Goodwin 2011, 293–95; Goldberg 2011)

In this chapter, I explicated what I mean when I speak about laypeople and their situational epistemic dependence on the more knowledgeable experts. Though they cannot directly assess the trustworthiness of expert testimony, they can use heuristic yet fallible ways to critically assess them. Such heuristics depend on social facts which are accessible to laypeople and justify them to make good judgments that are better than chance. While the philosophical debate mostly proposed passive heuristics, I will complement them with strategic considerations in this thesis.

3. The general discussion of the debate on heuristics for assessing experts

3.1 Assessing biases and incentive structures

¹⁷ I want to add that such an argument could take various routes: It could be understood temporally as in if this were true the news should have already reached me. Or by specific sources as in if that were true Miranda would also testify for it. Or finally, if that were true it would not just be you testifying for it. (see Goldberg 2011)

3.1.1 A conceptual analysis of biases and incentives

When laypeople try to assess which expert testimony they can trust, it is not enough to identify a person who has the best possible competence in the general discipline. As already mentioned in Chapter 2.2, we also need to be certain of the expert's honesty. Accordingly, Goldman proposes heuristics for identifying detrimental incentives in experts. When people have motives in competition to answer honestly, this reduces their credibility. For instance, Goldman cites a study on pharmaceutical funding where its authors argue that corporate third-party funding correlates substantially with more favourable conclusions for the latter. Such corrupting incentives are often instantiated by payment mechanisms, hope for prolonged procurements, or commission-taking. Following the simple rationale that when agents can benefit from being dishonest there is a risk that they will be, it is crucial for laypeople to know about such incentives. (see Crawford and Sobel 1982; Friedberg et al. 1999; A. I. Goldman 2001, 104)

Most scholars in the expert-identification debate concur with Goldman that corrupting incentives should be avoided by laypeople but rarely add anything more than additional examples of such problematic incentives. In this vein, philosophers like Goldman, Elizabeth Anderson, C.A.J. Coady or Jean Goodwin argue that corrupting incentives are a reason to, *ceteris paribus*, distrust putative experts independent of their competence. (see Anderson 2011, 147–49; C. A. J. Coady 1994, 204–5; A. I. Goldman 2001, 104; Goodwin 2011, 288–90)

This leaves open the question of how exactly laypeople should react to evidence of expert biases. Grundmann, for one, contends that testimony from biased sources should be disregarded entirely. I want to argue that this strict approach would have stern consequences as there exist no agents without the usual cognitive biases and hardly any whose interests do not affect their testimony at all. After all, experts have desires, moral values and are part of social groups just like laypeople. Empirical research, moreover, suggests that expert testimony is as prone to cognitive or ideological biases as that of laypeople as increased competence does not reduce such foibles. Ergo, treating detrimental incentives as a black-or-white matter would eliminate most testimony and figuratively throw out the baby with the bathwater. (see Angner 2006, 7–10; Grundmann 2024, 10–12; Morgan and Mellon 2011; Kahneman 2012)

Notably, Hein Duijf recently investigated the consequences of misaligned interests more closely by formally modelling their effects on the trustworthiness of experts.¹⁸ He explicitly analysed cases of gradually diverging incentives and how these relate to the competencies of experts and laypersons. While this is progress, there is still room for improvement as I will argue at the end of this chapter. Finally, Carlo Martini contends that the effect of expert biases on their testimony could only be mitigated by collective aggregation of diverse testimonies where different biases might partly cancel each other out and present a more robust perspective. I will consider such heuristics in Chapter 3.3. (see Duijf 2021; Martini 2014a, 9–11)

Apart from that, I want to argue that heuristics that assess the trustworthiness of putative experts by their biases and incentives are of special interest to my epistemic problem since they consist of social facts, which are directly accessible to laypeople. Even if finding and evaluating the magnitude of such incentives is not necessarily easy, not needing special expertise to access them reduces the laypeople's dependence on intermediaries. For instance, one does not need any insights into pharmacy or scientific operations to understand that funding from an agency that would benefit from a specific outcome gives scientists, who depend on funding, a reason to favour this outcome, and that this additional reason is not responsive to whether this outcome is the best result of proper research. Another example would be commission-taking often done by consultants who ought to use their special expertise to advise clients on making a suitable choice. Moreover, I contend that both the presence of interested funding and commission-taking are basic information in the sense that I explained in Chapter 2.2. Given the general difficulty of finding direct heuristics for laypeople, assessing biases and incentives is especially important. They can be the basis for finding meta-experts or selecting those whose agreement matters. For this reason, the insights in this section will reverberate in the following chapters. (see Collins and Evans 2009, 45–48; Collins and Weinle 2011, 407–11; A. I. Goldman 2001, 104–5)

¹⁸ In his paper, Duijf contends he is arguing against the 'widespread' belief that 'one should always trust expert, i.e. those who are more competent in a domain, regardless of misaligned incentives'. Duijf does not elaborate in which context this belief is supposed to be widespread. Contrarily, in the expert-identification debate multiple papers address how epistemically detrimental incentives can render competent experts untrustworthy: (see Anderson 2011, 147–49; C. A. J. Coady 1994, 124, 170, 185–86, 204–5; A. I. Goldman 2001, 104; Goodwin 2011, 288–90; Grundmann 2024, 10–12; Martini 2020, 117–19; Zagzebski 2012, 120–36). This is also the main insight of the classic cheap talk literature: (see Crawford and Sobel 1982; Farrell and Rabin 1996). I can only speculate that he is addressing statements in public media overcompensating against blanket denials of the epistemic authority of experts. Nevertheless, I think this is not a widespread belief in any discipline close to the expert-identification debate. (see Duijf 2021)

I want to continue by investigating the conditions that constitute such *epistemically detrimental* incentives. For that purpose, Kevin Elliott proposed three conditions while focusing solely on financial incentives. The ambiguity, the incentive, and the opportunity criterion. Ben Almassi, who emphasises that such incentives can also be driven by non-financial, idealistic motives, generalizes Elliott's conditions in this regard.¹⁹ Firstly, the existence of a conflict of interest depends on an ambiguity intrinsic to the testimony. This ambiguity criterion requires that the information is of a complex character that can lend itself to various interpretations. This is often the case with scientific data whose collection and interpretation can be guided by competing theories, whose choice is not obvious beyond doubt and open to reasonable disagreements. Propositions might also just be prone to misunderstandings due to their complex or abstract nature. (see Almassi 2017, 207–14; Elliott 2014)

I want to emphasize that this condition does not necessitate that information is ambiguous only that it can potentially be. Following the cheap talk literature, I discuss later in this chapter, I understand ambiguous information to be imprecise information, i.e. propositions of a higher generality. Consider these two examples of *actually* ambiguous information: 'A partial or full solar eclipse will happen within the next year' and 'a partial solar eclipse will happen on Thursday the 10th of July 2025 between 10am and 3pm CEST'. The first testimony is more ambiguous than the second. Accordingly, potential ambiguity could be seen as the range of actualities an information could denote a priori whereas actual ambiguous testimony depicts an interval of reality, such as 'the election will be held in the first quarter of the coming year' leaving open on which exact day. Due to the elusiveness of the factual world, most information will at least be a little bit ambiguous. In more technical terms ambiguity denotes a range of information where no probability distribution is being known. This also distinguishes it from uncertainty where neither the range nor the internal probability distribution is known. The reason such ambiguity can enable said problematic incentives is the theoretical opportunity, thus plausible deniability, it gives putative experts to choose an interpretation to their benefit. Conversely, 'basic information' as I characterised in Chapter 2.2, is a good example of testimonies that do not fulfil this condition and allow no room for imprecision. (see Almassi 2017, 207–14; Elliott 2014)

19 For instance, scientific experts might, intentionally or not, prefer interpretations of data that fit into their worldview or exaggerate evidence to increase their standing in the scientific community. Other, examples, would be (implicit) biases that move advisers to treat certain clients worse or take advantage of them out of disrespect. (see de Bruin 2021; Fidler and Wilcox 2021, chap. 2.2-2.5)

Secondly, Almassi's incentive criterion concerns the reason itself to not be fully honest, such as having financial stakes in particular propositions or other benefits. This is a necessary condition. (see Almassi 2017, 207–11; Resnik and Elliott 2013)

Thirdly, the opportunity criterion demands that individuals or institutions have the practical opportunity to slant the testimonies to their liking. While the ambiguity criterion refers to the theoretical possibility of interpreting evidence differently, this condition concerns the practical possibilities of agents to act on said theoretical possibilities, according to the detrimental incentives they have. In the scientific context, for example, mechanisms like peer review, open access data and mandatory disclosures of funding and interests can reduce such opportunities. (see Almassi 2017, 207–14; Elliott 2014)

I want to add, that from a lay perspective, the range of potential ambiguity in testified propositions is itself partly opaque. While general education might fairly calibrate laypeople's expectations of possible answers, the more precise range of potential ambiguity is not accessible without specialist insights. This turns some of the potential ambiguity into uncertainty.²⁰ Opportunities to bias testimony, on the other hand, are understandable without specialist expertise but might require extended information acquisitions on the institutionalization of expert advice in each context. This amounts to understanding the (exoteric) framework of how expertise is generated and mediated, how experts work together, and if or how mechanisms are structured to control their competence and honesty. More precisely this includes understanding which kind of testimony is to be expected by which experts, which bodies are supposed to check it and what this control means. For instance, the idea of peer review can be understood by laypeople and indicates reduced opportunities to bias testimonies. After all, to survive peer control the corruption necessary would need to be much more subtle or extensive. Hence, knowing about these increased obstacles can justify laypeople to judge testimony more trustworthy. Still, the epistemic benefits need to be weighed up against the effort of information transmission, especially considering the limited resources of the average layperson. (see Almassi 2017, 207–14; Lane 2014, 205–111)

Earlier I criticised that the expert-identification question has mostly remained rather vague about how (detrimental) incentives impact the trustworthiness of experts and what this means for laypeople. Recently, however, Duijf addressed this research gap and proposed a formal

²⁰ Asking additional testifiers might enable laypeople to narrow down the potential ambiguity of a proposition by observing the scope of differing testimonies.

model. At its basis are two probability variables; the likelihood of a layperson rightly identifying the suitable choice out of two options and the expert's likelihood to identify the layperson's suitable choice.²¹ Both the expert's and layperson's probability to succeed is between 50% and 100%, while it is conventionally assumed that the expert's likelihood is higher. Additionally, Duijf models a relative parameter, which denotes the alignment of the experts' preferences with those of the layperson by modifying the target of the expert's competence. This is Duijf's setup within which he compares several variations of competence probabilities and alignment sizes. Here, laypeople should defer to the expert iff the latter is more likely to identify the options that suit the layperson best. When interests are sufficiently misaligned, the expert's level of competence ceases to be an indicator of trustworthiness. Within his (two option) model, the interest alignment needs to be bigger than 0,5 so that higher competence justifies laypeople to defer. Depending on the competence difference between the two parties, different incentive alignment thresholds exist until when it would be rational for laypeople to believe the expert's testimony. This shows that testimony can still be advantageous to laypeople even if the expert is biased. (see Duijf 2021, 9295–9305)

Even though Duijf's model explicated what was just an intuition before, his approach is still lacking in some respects. Firstly, his probabilistic model treats suitable and non-suitable testimony as discrete outcomes, i.e. the testimony is, according to some probability, either perfectly informative or not at all. This again lacks nuance since intuitively laypeople often receive advice, which is neither fully dishonest nor completely truthful. Secondly, this model does not appreciate the strategic rationality of the involved agents since both competence and interest alignment are exogenous variables and thus Duijf misses how their heuristics and standings affect which indicators for competence are available and how detrimental incentives are being constituted. Ergo, this model remains in the passive tradition of philosophical expert-identification debate. (see Duijf 2021, 9295–9305)

There are nevertheless notable exceptions that treat the relationship of agents dynamically. For instance, Goodwin's account on reputation mechanisms, i.e., 'she would not risk...so I can trust', which involves the strategic anticipation of other agent's possible reactions. He, for example, characterises how public statements of experts can be seen as appeals to epistemic authority which put them in a position where they would lose reputation in the eyes of their audience if they were wrong, thus incentivising them to be not too overconfident. Also, the

21 Here, Duijf's research question slightly deviates from mine because I focus on cases where the layperson's competence cannot be significantly higher than chance. (see Duijf 2021)

critics of overvaluing consensus of experts, which I will return to in Chapter 3.3.1, concern themselves with the consequences of laypeople's actions on expert's incentives. (see Almassi 2017, 207–14; Goodwin 2011, 291–95; Intemann 2017; Martini 2014b)

To overcome the passive perspective in the expert-identification debate, I will now utilize game theory to model the relationship between experts and laypeople and to find out what properties affect the honesty of communication. More precisely, I will use the framework of cheap talk models, as they capture well the (potential) ambiguity in expert information, the effects of experts' (misaligned) incentives and the practical opportunities to bias they have, when there are no external constraints. This has the potential to show new ways how laypeople can improve their epistemic position within their limited resources. (see Almassi 2017, 207–14; Crawford and Sobel 1982; Goodwin 2011, 291–95; A. I. Goldman 2001, 104–5)

3.1.2 Strategic information transmission via 'cheap talk'

Cheap talk models are a specific type of game theoretic communication model, initially conceptualised in Crawford and Sobel's seminal paper "Strategic Information Transmission". The purpose of such game theoretic modelling is to predict and explicate how utility-maximizing agents would rationally interact, given their interests and other game-defining conditions. While real-life behaviour might involve a more complex set of motivations and constraints, abstract modelling can highlight underlying tendencies that guide real-life interactions. These insights can help us understand the dynamics between experts and laypeople beyond the passive information processing, presumed in the philosophical debate so far. (see Crawford and Sobel 1982; Farrell and Rabin 1996; Tadelis 2013, 357–65)

Typically, these games have two players; a sender that is informed about a state of the world and a receiver that could be made better off if she knew this state and could thereby act according to it. *Prima facie*, the sender and receiver have common interests, meaning that the sender also wants the receiver to act according to the true state. The game is structured so that, at first, 'nature' selects a sender type (the true state) within a commonly known random distribution of possible types.²² These are all possible states of nature, ergo the possible answers

22 This is also called the common-prior assumption. It is necessitated by Harsanyi's solution strategy for games with incomplete information that prior distributions of possible states of the world are known to all players. This is a simplification which ignores that the potential ambiguity is partially opaque to laypeople as I have argued earlier. For idealised modelling purposes, we might interpret this as the result of general school education calibrating laypeople's intuition about the range of reasonable answers. (see Tadelis 2013, 241–46)

to a question and might be thought of as being uniformly distributed between 0 and 1. In the second step, the sender privately learns ‘her’ actual type, the true answer to the question, and sends a message to the receiver including information on this type. When the receiver observes this message, she chooses an action that optimizes her utility according to this information. Finally, the sender and receiver each realise their payoffs according to the action of the receiver, each agent’s preferences for this action, and the true state of the world. Acting on the true state would always maximize the receiver’s utility, which is why she is motivated to ask an expert to begin with. The name ‘cheap talk’ refers to the fact that communication in such games is, contrary to signalling games, not directly payoff-relevant. Instead, outcomes depend on the action the receiver is being induced to choose due to this in itself costless message.²³ (see Crawford and Sobel 1982; Gibbons 1992, 210–19; Tadelis 2013, 357–65)

To clarify some terminological analogies: What economists call senders and receivers is a more general and technical version of talking about the relationship between experts and laypeople. The sender’s messages would be called testimonies (or more informally ‘advice’) in the philosophic debate. Sender types are possible answers to a question whereas the actual type of the sender, which she learns due to being an expert, is the right answer. The type- or state space is what I call the potential ambiguity of a testimony. Accordingly, cheap talk games capture the intuition of Almassi’s ambiguity criterion, which allows a ‘wiggle room’ in testifying for incentives to unfold their detrimental potential. (Perfect Bayesian) equilibria are predictions of the rational behaviour of agents in strategic interaction. In dynamic games of incomplete information several might exist, entailing that some predictions are more credible and intuitive than others apart from additional refinements that can be proposed. (see Almassi 2017, 207–14; Crawford and Sobel 1982; Tadelis 2013, 357–65)

Analysing the model, if the sender decides to message the true state and if the receiver believes that the sender is honest, then the receiver’s best response is to assume the reported state since acting on the true state is always optimal for her. The sender, anticipating that the receiver would react like this, would have no reason to deviate from her truth-telling strategy

23 This structure makes it a dynamic game of incomplete information, as agents are playing in consecutive turns and receivers are ignorant about the sender’s actual type, though not about the range of potential types and their prior likelihood. Also, each player’s possible actions, rationality and rules of the game are common knowledge. The solution concept that predicts agent interaction with sequential rationality, is the perfect Bayesian Nash equilibrium. In short, the latter requires that the conjectured strategy of every player is a best response regarding their beliefs about the other player’s action, meaning that no player could benefit by deviating from this strategy. Furthermore, the beliefs regarding each action of the other player need to be consistent with the player’s strategy and credible regarding possible actions beyond the conjectured strategy. (see Gibbons 1992, 173–83; Tadelis 2013, 303–14)

as the receiver's action is also her preferred choice. The sender does this by conditioning the message she sends on the information she received about her type. After all, the sender has no reason to deviate since they share the same preferences, and any differing message would make her worse off in comparison. In this case, the sender's message is fully truthful and informative and both parties realize the best possible payoffs. Such a situation where neither agent has a reason to deviate, and the information is communicated without loss is called a fully revealing (perfect Bayesian Nash-) equilibrium. The first trivial insight is that communication can alleviate information asymmetries or can, at least, improve the situation of the less informed. (Crawford and Sobel 1982; Gibbons 1992, 210–19; Farrell and Rabin 1996; Tadelis 2013, 361–65)

While this would be the most informative equilibrium other less informative ones could be constructed as well. Imagine the sender chooses messages independently of the information she has about her type. In equilibrium, the receiver knows the sender's strategy, i.e., she knows that the sender's message is not correlated with the sender's signal. Consequently, it is optimal for her to ignore the sender's message and base her action on her prior probabilities only. Ignorance of the sender's message by the receiver in turn comforts the sender in her choice to send messages that are independent of her information. No player has a reason to deviate from this strategy, beliefs are consistent, and no information is transmitted. Such an equilibrium is called a babble equilibrium, just as if the sender would speak unintelligibly.²⁴ (see Crawford and Sobel 1982; Farrell and Rabin 1996, 107–10; Sobel 2013, 317; Tadelis 2013, 361–65)

The interesting part about cheap talk games, however, is what happens when the sender's preferences are not perfectly aligned with those of the receiver. This is modelled by including a bias variable, which slants the action a sender would prefer the receiver to choose away from the receiver's optimal point. Now we reach a situation where the sender still knows the true information (by being an expert) but might prefer that the receiver deviates from the true state.

²⁴ Both the babbling case and other more informative equilibria fulfil the basic conditions for Perfect Bayesian Nash Equilibria which opens the question, of which of them truly predicts agent behaviour. There are various proposals in the literature to best select equilibria in such scenarios. Joseph Farrell and Matthew Rabin, for example, argue that babbling equilibria simply are counter-intuitive modelling artefacts, since players share a natural language for coordination and a common interest to abide by more informative equilibria. This has been coined the rich-language assumption, where cheap talk may be ineffective due to incredulity but not incomprehension. Furthermore, Archishman Chakraborty and Rick Harbaugh show that a sufficient condition for both sender and receiver is to weakly prefer informative equilibria over the babbling one if they have quasi-convex preferences. Another way to manage the equilibrium selection problem is to just choose the most informative equilibrium. While properly justifying this might require a more detailed discussion, I contend that, for my research question, choosing the most informative equilibrium is an inconsequential shortcut. (see Chakraborty and Harbaugh 2010; Crawford and Sobel 1982; Farrell and Rabin 1996, 107–10; Gibbons 1992, 210–19; Sobel 2013, 308–11, 313–17; Tadelis 2013, 375–365)

The bias is common knowledge, just as each agent's preferences are. For example, a sender might prefer a receiver to subscribe to a proposed tariff one level higher than what would be optimal for the receiver, because then the sender gains a higher commission, while the receiver is still satisfied enough to stay a customer. (see Crawford and Sobel 1982; Gibbons 1992, 210–19; Tadelis 2013, 361–65)

I now want to compare the scenario where the sender has a bias with the description earlier when both agents shared the same preferences: If the sender decides to perfectly correlate her message with the true state, then, in equilibrium, the receiver knows this, and it is the best response for her to optimize according to this state. When the sender anticipates this reaction, however, she has a reason to deviate. Knowing that, given the receiver's belief, she would act according to the message she sends her, the sender would be better off, if she modified the message by her bias. Perfect revelation is thus not an equilibrium. (Babbling, however, remains one.) When the receiver believes that any sender type would send the same message independent of the true state her best response would be to maximize the entire type space given the prior probabilities, hence using no information. However, if the sender's bias is not too large, she can use an alternative strategy both would prefer over the latter outcome. (see Crawford and Sobel 1982; Gibbons 1992, 210–19; Tadelis 2013, 361–65)

The sender can partition her type space into two intervals divided by a certain threshold type. Her strategy would, then, be to choose one message if the true type is below this threshold, and another message if it is above this point. In equilibrium, the receiver knows that the sender plays such a threshold strategy, and she learns in which of the two intervals the true state is. Her best response is thus to choose an action given the weighted average of the reported interval of possible sender types (given the prior probabilities which are uniformly distributed anyway). For a small enough bias, there always exists a threshold such that sender types below and above the threshold each prefer the receiver's action that optimizes on their interval over other actions, i.e., a threshold such that no sender type has a reason to deviate. As neither of them can profit from deviating and their beliefs are consistent with these strategies, this is a perfect Bayesian Nash equilibrium. Here, notwithstanding the bias of the sender, some information has been transmitted. The receiver learned that the true type lies somewhere within one or the other interval, which is more precise than what she knew at the beginning, namely that the true type lies within the entire type space. (see Crawford and Sobel 1982; Gibbons 1992, 210–19; Tadelis 2013, 361–65)

This result crucially depends on the size of the bias. If it was too big, the sender's interests would differ from the receiver so much that any sender type would benefit from messaging a higher type regardless of the true state, conversely, the true state cannot be low enough to induce the sender to message anything else. Consequently, partially revealing equilibria are only possible below a certain threshold of bias size and this threshold is relative to the type space (and indirectly to the rationalizable message and action space).²⁵ Crawford and Sobel furthermore show that small biases allow strategically stable partitions with more than two elements and the smaller the bias gets the finer the partitions can be, i.e., the more information can be transmitted. (see Crawford and Sobel 1982; Gibbons 1992, 210–19; Tadelis 2013, 361–65)

Unfortunately, this misalignment in preferences entails that whatever the true state of nature is, if the bias is not 0, the sender would always gain by deviating from truth-telling and would report an imprecise message instead. Therefore, full truth-telling is not an equilibrium strategy of the sender, i.e., there is no fully revealing equilibrium in situations where the sender has some bias.²⁶ (see Crawford and Sobel 1982; Farrell and Rabin 1996; Gibbons 1992, 210–19; Tadelis 2013, 361–65)

3.1.3 A game theory informed analysis of incentive structures

The most central insight the standard cheap talk literature lends formal support to is that honest testimony transmission crucially depends on (fully) aligned interests between the expert (sender) and layperson (receiver). Taking this seriously prevents us from mistakenly asking for disinterested experts. Colloquially speaking, sharing interests might be associated with subjectivity, but in a scenario of advice, the preferences of the client are the proper object of concern. In a similar vein, C.A.J Coady writes that disinterestedness can entail epistemic vices like lack of attention and therefore accuracy. Duijf further adds that misaligned incentives do not need to be born out of malice. Incidental misunderstandings can move an expert to give biased advice. Also appeals to a generalized public good or paternalism for the individual, could

²⁵ Crawford and Sobel demonstrated in their standard case that the bias variable needs to be smaller than $\frac{1}{4}$ relative to the type for there to exist a partially revealing equilibrium. (see Crawford and Sobel 1982, 1440–44)

²⁶ While cheap talk models clearly analyse the general problem with corrupting incentives, they are not easily applied to scenarios where the incentive cannot be easily modelled as a (linear) variable. Such cases could be non-financial or idealistically driven incentives. Almassi worries in this regard: “(...) that non-financial COIs are neglected not because they are unimportant but rather because of the constraints of operationalization that publication requirements and codes of professional ethics demand.” (2017, 209)

motivate experts to impose interests onto laypeople who they do not share. The *source* of interest misalignment has, however, no impact on the modelling setup. (see C. A. J. Coady 1994, 203–5; Crawford and Sobel 1982; Duijf 2021, 9291–93; Farrell and Rabin 1996)

Contrary to Grundmann’s strictness, the cheap talk analysis shows how biased expert testimony can still have some epistemic value when the bias is small enough. I suspect that the passive perspective of major parts of the philosophical debate leads to the short-winded proposals towards laypeople to simply avoid biased experts or vaguely to ‘reduce trust’. Already Duijf showed that biased experts can still be trustworthy enough if the bias is small enough. (see Duijf 2021; A. I. Goldman 2001, 104–5; Grundmann 2024, 10–12; Martini 2020, 117–19; Scholz 2009, 194–202)

The cheap talk literature, however, explicated in a more nuanced fashion how biased expert preferences lead to imprecise testimony transmission. I transpose this game theoretic rationale into the following de-idealised real-world dynamic: A biased expert knows the precise truth but, *prima facie*, has a reason to communicate a precise but slanted testimony instead. Rational laypeople (perfectly) anticipate the bias but do not know the true point from which the expert deviated. Instead of risking misleadingly discounting this testimony, it is rational for them to believe only in the general gist of the expert testimony. Depending on the size of the bias they should believe in the precision of testimony only as far as the bias could negate it. Consequently, laypeople are better informed than before the expert’s testimony, but still only know a less precise more ambiguous account of the world since the expert cannot be trusted for more precision. The latter also prefers that laypeople only know these more ambiguous accounts rather than fully enlighten them. In this vein, for example, reducing trust could mean believing a personal-finance expert insofar that it is sensible to buy household insurance but not which specific product (which this expert might get a commission for).²⁷ (see Crawford and Sobel 1982; Farrell and Rabin 1996; A. I. Goldman 2001)

A cheap-talk-informed understanding of biased expertise further appreciates the strategic relationship and impact of Almassi’s (potential) ambiguity and opportunity criteria. This also shows that there is effectively no chance of a biased expert to testify truthfully because this is not in the latter’s interest. Not simply disregarding biased advice is especially sensible, since in

²⁷ Note, that I am not claiming this represents Crawford and Sobel’s cheap talk model. After all, there the sender anticipating the receiver’s anticipation of her bias, messages an ambiguous message to prevent the receiver from being able to mathematically deduct the bias, which consequently is an equilibrium strategy. Instead, this is an interpretation of ‘normal language’, which connects the insights of the model with the philosophical and sociological debates. (see Crawford and Sobel 1982; Farrell and Rabin 1996; A. I. Goldman 2001)

messy real-world scenarios, smaller detrimental incentives may be found frequently. Just think about sales talks or political consulting. Considering the insights from the cheap talk models, we should start a more nuanced discussion on how laypeople could utilize imprecise information best. (see Duijf 2021; A. I. Goldman 2001, 104–5; Grundmann 2024, 10–12; Martini 2020, 117–19; Scholz 2009, 194–202)

Note that the cheap talk model discussed so far, treated the expert as being perfectly informed. If the signal she learns involves intrinsic uncertainty, the expert's competence is impeded and only a reduced amount of knowledge can be transmitted. In Chapter 5.3.2, I will discuss a game theoretic model with multiple senders where this possibility becomes even more consequential. Unfortunately, expertises often contain at least some intrinsic uncertainty while others involve even comparatively more. In the applied part of this thesis, I will investigate this issue in the context of personal-finance expertise. (see Ambrus and Lu 2014; Crawford and Sobel 1982; Lane 2014)

A further assumption cheap talk models make to analyse the strategic communication is, that the size of the bias is common knowledge. In the model, the layperson's knowledge about this size does not entail that she can discount it, but it is necessary for her to partly disincentivise, i.e. decline, a misleadingly precise message and accept an ambiguous testimony, whose imprecision depends on this size. In a real-world context, the laypeople's assessment of the size of the bias is the result of her heuristic judgments and thus fallible. After all, cheap talk models abstract away from the laypeople's inability to perfectly assess the potential ambiguity of testimony. In addition to the limited means of laypeople I described in Chapter 2.1, I think we should not assume that laypeople could perfectly assess the magnitude of an expert's bias. Instead, it is a part of the heuristic task of laypeople to detect these biases and approximate their true size by using the criteria I have discussed so far and will continue to in the rest of this chapter. (see Almassi 2017, 207–14; Crawford and Sobel 1982; Farrell and Rabin 1996)

Considering the opportunity criterion, I want to propose that the incentives of experts can be reduced to the degree laypeople are able to (a) detect epistemically detrimental incentives and (b) punish experts who indulge in these incentives. The first part (a) is self-explanatory, if laypeople ought to restrain expert misbehaviour they need to know of its existence. Punishing experts (b) could take the form of filing a lawsuit to achieve a fine for fraudulent behaviour, or avoiding future business relations to incur income loss, damaging their reputation, and so on. The more efficient laypeople are in detecting and punishing detrimental biases, the less lucrative are these incentives as putative experts can anticipate such consequences. The efficiency

crucially depends on the expert's sensitivity to punishment and the layperson's bargaining power. (see Almassi 2017, 207–14; Crawford and Sobel 1982; Goodwin 2011, 293–95; Farrell and Rabin 1996; Sobel 2013)

It is also the purpose of the strategic-dynamic heuristics, which I am emphasising in this thesis, to detect and punish corrupting incentives. A central mechanism relevant here is reputation, which is addressed by Goodwin in his normative pragmatic account of the appeal to expert authority. Colloquially speaking, it follows the idea of 'he would not risk...so I can trust', which Goodwin deems a widespread principle in our society alleviating information asymmetries. Note, however, that not every putative expert is equally sensitive to reputation loss. For example, some might be in a monopoly position which laypeople cannot avoid. Alternatively, potential ambiguity or practical opportunities could be serious enough so that bad outcomes of advice cannot be effectively attributed to expert failure. To no surprise, then, putative experts have an interest in avoiding responsibilities and maintaining plausible deniability. (see Goodwin 2011, 293–95)

More generally, laypeople could improve their ability to detect biases by referring to meta-experts. Goldman describes these as agents who assess the competence and honesty of other experts. They examine candidates via their expertise and certify them for laypeople, which often happens in an institutionalised context.²⁸ In practice, this could take the form of educators who examine their students and issue a degree. Such degrees or labels conventionally make the relayed trust visible to laypeople. Another example would be certifications for proper conduct and competence, which professionals can often receive from a public meta-expert agency or a cachet from a private standard board. Of course, for such meta-experts to be able to relay trust, they first need to be trusted themselves. This renders the deference to meta-experts an indirect heuristic judgment, which necessitates laypeople to use other direct heuristics first to justify their trust in them. Otherwise, the justification would be circular or involve an infinite regress. If laypeople, however, have enough reasons to trust these meta-experts, the latter can effectively help extend the circle of trustworthy experts in ways laypeople could not do themselves. (see C. A. J. Coady 1994, 282–84; A. I. Goldman 2001, 97)

28 This kind of interaction between experts and laypeople can also be modelled by a game theoretic communication model. In 1973, Michael Spence first described a signalling model where a potential employee signals her future employer to be of high competence by having a university degree which warrants a time and energy sacrifice which can only be rational for high competency types since for such endeavours are easier than for low competency types. Ergo, without knowing the competence, an employer can rationally infer this information by observing this credential. (see Gibbons 1992, 183–89; Spence 1973; Tadelis 2013, 318–23, 332–35)

Finally, I want to argue that the legal framework of a jurisdiction can substantially impact the trustworthiness of putative experts from a lay perspective. After all, it could constrain problematic behaviours of each party, assign punishments for violations and regulate which and how institutions control said practices and infringements. Such frameworks can incentivise better behaviour to the degree they distinguish well between desirable and deceitful practices, assign suitably motivating punishments, and establish mechanisms that are effective in recognising these behaviours. The effective threat of punishments would reduce the scope for biased testimonies since too ambiguous testimonies would not be beneficial anymore. In turn, laypeople can anticipate that such testimonies will not be highly biased and trust more precise information. Nevertheless, laws cannot fully eliminate bad advice and need to focus on punishing clear violations, i.e. cases of proper fraud, since control institutions are unlikely to have the resources to engage in (expert) negotiations on too specific matters. In simplified terms, legal frameworks in (quasi-) democratic contexts typically are the result of negotiations between stakeholders and will to different degrees represent different interests. Accordingly, different legal frameworks in different jurisdictions will vary in how effectively they license trust for laypeople. This aspect has not received explicit attention in the expert-identification literature so far. I will argue in Chapter 5.1.3, that this is likely caused by the debate's implicit focus on (natural) scientific expertises, whereas consumer protection law is more relevant in practical, consulting expertises like personal finance.

Relatedly, scholars like Grundmann and Anderson emphasise how scientific institutions are designed to objectify knowledge and disincentive fraud, or malpractice. In this vein, social epistemologists of science debate the effectiveness of double-blinded Random Control Trials (RCTs), peer review or reputation mechanisms. (see Anderson 2011; Longino 2019; Grundmann 2024, 15–17)

In this chapter, I extensively discussed how detrimental incentives can impact the trustworthiness of (putative) experts. Using game theory, I showed how misalignments of the interests of laypeople and experts will degrade the epistemic value of the latter's testimony. If the expert's bias is, however, small enough a layperson can still extract helpful information by only believing in the more general trust of the message proportional to the size of the bias. In this respect, I have complemented (and partly corrected) the philosophical debate on the expert-identification question by characterising a more nuanced understanding of detrimental incentive structures. This bias and its size ought to be heuristically approximated by the layperson by following the criteria of Almassi. Accordingly, laypeople should follow a general heuristic of

‘detecting biases and assessing their bargaining power in disincentivising them’. They can do this by appraising the potential ambiguity of the expert testimony and by assessing the practical opportunities experts have to act on their biases along with how laypeople’s bargaining power or present institutional sanctioning mechanisms could affect the honesty of expert testimony. (see Almassi 2017)

3.2 Directly assessing the past track record of putative experts

I will now address Goldman’s most original contribution to the expert-identification question, which he also deems the most helpful method. Although laypeople normally cannot directly assess expert testimony, as I discussed in Chapter 2.2, Goldman noticed an exception to this rule. When a putative expert, for instance, makes a prediction or repairs a system, laypeople might at some point observe its success condition realising. Observe his example: (see A. I. Goldman 2001, 106–8)

For example, consider the statement, ‘There will be an eclipse of the sun on April 22, 2130, in Santa Fe, New Mexico.’ Relative to the present epistemic standpoint, i.e., the standpoint of people living in the year 2000, this is an esoteric statement. Ordinary people in the year 2000 will not be able to answer this question correctly, except by guessing. On the other hand, on the very day in question, April 22, 2130, ordinary people on the street in Santa Fe, New Mexico will easily be able to answer the question correctly. (A. I. Goldman 2001, 106)

For this purpose, amongst others, Goldman introduces the distinction of esoteric and exoteric elements of expert testimony, where ‘esoteric’ denotes non-accessibility to laypeople and exoteric the opposite. In his (astronomical) prediction example, former esoteric testimony turns exoteric at some point in time and position on Earth. When laypeople observe a prediction realising, they have good reasons, *ceteris paribus*, to infer that the putative expert must have had some special insight earlier, when its occurrence was not exoteric. Goldman also gives other examples of the same principle, such as scenarios where ‘systems’ are being repaired like cars or human bodies. Furthermore, when someone builds a machine, whose success condition is salient to laypeople, for instance by returning it to its proper function, laypeople can

heuristically infer that this ‘treatment’ might have been the product of expert competence. I will call this the ‘salient-success’ heuristic in my thesis.²⁹ (see A. I. Goldman 2001, 106–8)

Goldman is aware that such assessment opportunities are not available in every context and expertise. He argues, however, that this possibility to directly assess expert testimony allows us to justifiably trust experts, who in turn can certify colleagues whose expertise is not directly accessible by laypeople, widening said justified trust in an expert community. Unfortunately, Goldman is not clear which kind of ‘salient success’ can justify trust in an expert by roughly how much. Here, I think it is helpful to take a closer look at his examples and analyse the intensity of the heuristic evidence involved. Referring to Stegenga, the degree of intensity of evidence depicts the amount of support one observation can give to the conclusion. Successfully building a rocket or predicting a solar eclipse to the exact day years into the future, are golden-event observations that give a lot of support for a very rare event. Accordingly, it might suffice to succeed in a few such endeavours for laypeople to infer expert competence justifiably heuristically. (see A. I. Goldman 2001, 106–8; Stegenga 2009, 654–56)

Goldman and other philosophers directly responding to this heuristic, however, failed to explicate that the same is not true for weather forecasts, car repair, medical treatments, and the like. Observed success in the latter context gives lower intensity (heuristic) evidence for more frequently occurring events. Connecting this debate to the discussions of Philip Tetlock and Erik Angner, laypeople need a wider picture of the general accuracy of putative experts in low-intensity cases, i.e. a record of being right compared to being wrong plus the reported confidence, to make more reliable inferences. After all, being incidentally right is far more likely for such predictions, especially if one just makes enough predictions. Especially, when testimony is potentially ambiguous, i.e., can be interpreted in multiple ways so that it practically encompasses a broad range of future possibilities, an actualisation of one of these possibilities might be quite likely. Certainly, transparently making broad predictions is unproblematic since laypeople can benefit also from more cautious and less competence-demanding testimonies. Intentionally ambiguous testimony, however, allows putative experts to hide intrinsic

²⁹ Just as in footnote 16 in Chapter 2.2, I want to note Paul Milgrom and John Robert’s game theoretic models that analyse information transmission between an expert and a layperson, where the testimony is verifiable for the latter after receiving it. The gist of their work is that, if testified information is post-hoc verifiable, then the sender cannot benefit from sending biased messages, since the truth would be revealed anyway. There is however a way for senders to improve their position by strategically withholding information, i.e., experts might prefer testifying the part of the information that benefits them and leaving out inconvenient aspects. Applying these insights to this chapter, experts could have reasons to avoid testimonies which become exoteric at some point by giving ambiguous testimony, for example, by omitting probabilities in forecasts or using vague time or intensity quantifiers. (see P. R. Milgrom 1981; P. Milgrom and Roberts 1986; Sobel 2013, 307–8, 327–28)

uncertainty and might feign more competence than it involves. Tetlock and Angner, for example, show that political forecasters do not face proper sanctions for inaccurate and ambiguous predictions. I will further discuss this issue when I speak about forecasting in personal finance in Chapter 4.3. (see Angner 2006, 10–14; A. I. Goldman 2021, 89–91; 2001, 106–8; Grundmann 2024, 12–13; Stegenga 2009, 654–56; Tetlock 2017, xxxv–xxxvi)

In his 2021 paper, Goldman noticed that cases where testimony becomes as unambiguously exoteric as in golden-event observations are rather scarce. Experts often make predictions about for laypeople unobservable phenomena, such as theoretical predictions of quantum particles which are tested in quantum particle accelerators. Nevertheless, he writes: (see A. I. Goldman 2021, 89–91; Grundmann 2024, 12–13)

To take simpler examples, laypersons can sometimes successfully identify cases in which their doctor, their mechanic, their investment adviser, etc. gives them just the sort of treatment, advice, etc. that palpably succeeds in generating sought-after results. (A. I. Goldman 2021, 90–91)

Even though Goldman presents this ‘salient-success’ heuristic as only occasionally applicable, his examples still imply an overtly optimistic attitude towards laypeople’s ability to differentiate competence-driven successes from incidental ones. While I agree that observing a healed body or a working car after treatment gives us, *ceteris paribus*, heuristic reasons to believe in the expertise of the agent, the scope of this heuristics is limited. After all, the causal structures involved in complex systems such as cars and human bodies, are difficult to navigate for laypeople, since such cases tend to involve more intrinsic uncertainty as I will argue in Chapter 4.2. For instance, it is not clear whether the putative expert only applied a temporary fix, whether she sold a treatment that was not necessary or whether the outcome was not driven by the putative expert (the body healed itself over time). Car repair is a notorious example of a credence good, further discussed by Uwe Dulleck and Rudolf Kerschbamer. The situation for investment advice is even more peculiar, which is why I will analyse it in the applied part of my paper and argue that the salient-success heuristic is useless in this context. (see Dulleck and Kerschbamer 2006; A. I. Goldman 2021, 89–91; Grundmann 2024, 12–13)

Finally, this prompts the question of how laypeople could learn about the intensity of such heuristic evidence for success. I fear this is a difficult task a layperson could only approximate by having some minimal understanding of the character of the expertise, i.e. by learning how easy it is to succeed without the expertise and how frequently involved events happen in general. Such a minimal understanding could be gauged by comparing putative expert

testimony or learned by meta-expert testimonies, which themselves need to be trusted. I will now discuss the heuristics of the former kind in the following chapter. (see A. I. Goldman 2021, 89–91; 2001, 106–8; Stegenga 2009, 654–56)

3.3 Comparing testimonies of multiple experts

3.3.1 The epistemic value of agreement and consensus

The heuristics so far addressed putative experts one by one, though they naturally come in groups. It is questionable whether highly specialised knowledge could be cultivated by individuals disconnected from peers. On the contrary, scholars of social epistemology and the sociology of experts and expertise (SEE) ground their analysis in the intricate *social* interactions necessary for knowledge generation and transmission. Remember, for example, C.A.J Coady's argumentation in Chapter 2.2 regarding testimonies as a basic source of knowledge or the discussion of how laypeople could heuristically estimate an expert's trustworthiness by judging their social position in our society. It seems intuitive to take advantage of the collaborative nature of expertise and compare the testimonies of multiple experts. Is it not common wisdom many of us learn in school to consult different sources to receive more robust, complete, or less biased information? Perhaps propositions receive majority approval for good reasons? (see C. A. J. Coady 1994; Collins and Evans 2009; Collins and Weinel 2011; A. Goldman and O'Connor 2021; A. I. Goldman 2001)

Indeed, heuristics addressing this intuition have received a lot, if not the most, attention in the philosophical debate following Goldman. While the majority seems to attach a central role to comparing the agreement of experts, there is disagreement regarding 'how' to count it. Nevertheless, there are also critical voices that warn of overestimating the value of consensus. Hence, I will begin by generally discussing the epistemic value of agreement and consensus before I move on to discuss lay heuristics more specifically. (see A. I. Goldman 2001; Miller 2020)

David Coady, for instance, contends that heuristics of 'counting agreement' are the only ones not presupposing (meta-)expertise without substantiating his judgment. I decisively disagree, since such heuristics already assume that we have some understanding of who is eligible to be compared. Otherwise, we risk diluting proper consensus assessments with arbitrary accounts from charlatans, reducing the meaningfulness of such heuristics substantially. This makes this endeavour one of 'indirect' heuristic judgments like the

assessment of certifications by meta-experts. First, laypeople need to use their direct heuristic means to acquire a finer selection of putative experts before they can meaningfully use these aggregational methods. This also means that ‘counting agreement’ cannot be any less demanding than ‘the assessment of incentive structures’, which David Coady instead deems as requiring some expertise. I, however, have argued in Chapters 2.2 and 3.1 that the latter heuristics only require the understanding of social facts, for which laypeople are properly equipped. In contrast to David Coady, I am arguing that laypeople can, for instance, assess the incentive structure in a discipline and the biases of individual putative experts to decide which of their testimonies is trustworthy enough to be considered in aggregation. (see D. Coady 2006, 75–76; Grundmann 2024, 15–17)

In his seminal 2001 paper, Goldman argues that laypeople should count the agreement a proposition has within the relevant expert community and weigh it up against the agreement towards alternative positions. He resorts to the intuition that each putative expert’s testimony gives epistemic weight towards a proposition just as additional eyewitnesses increase the credibility of the perception they support. *Prima facie*, the greater number of agreements should be one desideratum amongst other heuristics for choosing expert testimonies, though he does not spend much time motivating this intuition. (see A. I. Goldman 2001, 97–99)

Indeed, this intuition is not shared by Miriam Solomon who argued that the true indicator for empirical knowledge is empirical research and not social deliberation. Promoting consensus would be superfluous, since either they are coerced prematurely, or they already instantiated themselves naturally, rendering the endeavour redundant. In a similar vein, Kristen Intemann warns that the pressure to form a consensus might silence disagreeing voices and thereby valuable alternative explanations. Dissent is an essential part of (scientific) expertise and a prerequisite for the critical testing of theories. Relatedly, Boaz Miller reminds us that just as individual experts can be wrong, neither is the majority view exempted from this possibility. Both Intemann and Martini add that already laypeople’s expectation that agreement is a sign of quality could set unrealistic standards for (scientific) experts and would incentivise them to be conservative in their research or overtly seek peer validation. After all, expert’s funding and esteem often depend on their public image. Moreover, Solomon and Miller argue that an expert community with many members does not easily agree on a consensus, since there can be mutual misunderstandings, personal rivalries, and other communication problems. Conversely, different groups of scientists share non-epistemic goals, biases, and interests, allowing also misguided consensus to happen. Experts, for instance, could fake a consensus to protect the

reputation of their discipline. Instead, a consensus might primarily indicate trust between experts rather than a trustworthy position. (see Intemann 2017, 198–201; Martini 2014b; Miller 2020, 230–33; 2013, 1293–98; Solomon 2015; 2007, 117–37)

The concerns of Martini and Miller but especially of Solomon and Intemann bring attention to the fact that agreement within an expert community is neither the natural goal nor normative aim of science (or related expertises). Note, however, that the philosophers of the expert-identification debate are not addressing the norms of the scientific process but fallible heuristics for laypeople that justify their trust in a difficult epistemic context. While overvaluing agreements between experts could mislead laypeople and inhibit the progress of a discipline, this does not entail that agreement cannot have some heuristic value. After all, Martini raises also pragmatic arguments for the epistemic value of agreement: Politicians and civil servants at some point require non-contradicting recommendations to inform policymaking. Collins and Evans add that the speed of politics is faster than that of science (and associated expertises) and thus experts often are pressed to make pre-emptive authoritative decisions. This argument also directly translates to the needs of responsible laypeople. Since agreement between expert testimonies and their indicators are social facts, they are accessible to laypeople and of indispensable heuristic value. (see Collins and Evans 2009, 2–9; Martini 2014b, 115–21)

The critique notwithstanding, heuristics that ‘count agreement’ are considered sensible by most scholars of the expert-identification debate, like, for instance, Anderson, Martini, Miller, David Coady, C.A.J. Coady, Grundmann, Collins and Weinle, and Goldman himself. I concur that the more interesting question is rather how laypeople could identify whether a consensus is the result of an alignment of technical expert reasoning or an agreement due to non-epistemic social factors. Miller addresses this issue by conceptualising four possible kinds of consensus deriving from a matrix of ‘originating in a justified/unjustified process’ and the view being ‘right/ wrong’. Firstly, a consensus can be *non-cognitive*. The function of these is to establish a useful convention like how to cite papers in a journal, which enables cooperation between scientific experts. Secondly, consensus can be *vertically lucky*, when the position is accidentally right, without originating in a justified process. Thirdly, a consensus can be *epistemically unfortunate* by deriving from a justified process but being nevertheless wrong, due to systematic biases or deliberate deviation of experts. If a consensus falls into none of these categories, it is likely *knowledge-based* and thus epistemically valuable for lay heuristics. (see Anderson 2011, 147–49; C. A. J. Coady 1994, 283–85; D. Coady 2006, 68–75; A. I.

Goldman 2001, 79–104; Grundmann 2024, 15–17; Intemann 2017; Martini 2020, 117–19; 2014b; Miller 2020, 231–36; 2013, 1300–1314; Solomon 2015; 2007, 117–37)

3.3.2 The robustness of expert testimonies

In accordance with the expert-identification debate, I now move on to discuss the attempt to identify criteria which separate epistemically valuable agreements to a testimony from those that should be neglected. Starting with Goldman's approach, he specifies that, *ceteris paribus*, testimonies need to be 'independent' of each other, so that their agreement justifies the layperson to increase trust in the agreed-on proposition. He stipulates what Jennifer Lackey calls the 'belief-independence-thesis'³⁰: (see A. I. Goldman 2001, 97–104; Lackey 2013)

If two or more opinion-holders are totally non-independent of one another, and if the subject knows or is justified in believing this, then the subject's opinion should not be swayed - even a little - by more than one of these opinion-holders.³¹ (A. I. Goldman 2001, 99)

Goldman uses a Bayesian analysis of new testimonial evidence updating the prior beliefs of interlocutors to explicate the difference between testimonial independence and dependence and why the latter negates additional epistemic weight. An additional testifier's belief is independent of the original testifier if they do not share the same causal routes from evidence to belief. Following Goldman, this, *ceteris paribus*, increases the likelihood of the additional testifier agreeing with the original testifier if the latter is right rather than when the latter is wrong. In Bayesian terms, additional testifiers need some kind of 'new' evidence that is not already included in the prior beliefs, like access to independent eyewitnesses, experiments, or processes of reasoning towards the truth. This can also be a gradual matter, where testimony is partially dependent on the original testifier's belief but not fully, leaving open room for correction. (see A. I. Goldman 2001, 97–104)

To illustrate this point, Goldman uses the example of a guru whose followers uncritically copy his sayings. The epistemic weight one guru adds should be equivalent to the weight a guru plus his uncritical followers add. Thus, laypeople should not count the agreement of non-discriminating reflectors as extra heuristic evidence for a position. In a similar context, Ludwig

30 David Coady calls it the non-independence principle, but I will use Lackey's formulation in this thesis. (see D. Coady 2006; Lackey 2013)

31 Goldman chose the double negative 'non-independence', though I see no reason to not just call it dependence. (see A. I. Goldman 2001, 97–104)

Wittgenstein compared epistemic support from non-independent sources to buying several copies of the morning paper to assure oneself that what it said was true. (see D. Coady 2006, 68; A. I. Goldman 2001, 97–104; Wittgenstein 1998, para. 265)

Besides scholars endorsing or refining Goldman’s heuristic of ‘comparing the agreement of experts on a testimony’, Miller parallelly proposes three ‘indicators’ which ought to be present to make it more likely that a consensus is knowledge-based: social calibration, apparent consilience of evidence and social diversity. Participants in a consensus are socially calibrated if they share the same formalisms, evidential standards, and ontological schemes. Apparent consilience of evidence, further, is the result of diverse methods or data that seem to converge. Finally, a consensus is socially diverse if expert testifiers with various social identities converge on a position. This should reduce the impact of idiosyncratic biases prominent in distinct social groups since the testimony is also being supported by experts not sharing the same bias which reduces the chance that the testimony has been (overtly) driven by this bias. While Miller calls these criteria ‘indicators’, I think this term is misleading as they are relevant conditions for knowledge-based consensus, but laypeople cannot directly use them to check testimony. Instead, they first need to make other observations that *indicate* the fulfilment of these conditions.³² (see A. I. Goldman 2001, 97–104; Miller 2013, 1300–1314)

Having these contributions in mind, I want to argue that these all (implicitly) are robustness-style arguments. While scholars, so far, remained rather vague concerning the underlying argument for ‘counting agreement’ and emphasised its intuitiveness instead, I think it is helpful to introduce Jacob Stegenga’s thorough analysis of robustness. Miller did address this concept as well yet merely to justify the epistemic value of expert agreement not to better inform the related heuristics, as I intend to do. (see Miller 2013, 1308–12; Stegenga 2009)

32 There is just another rationale closely associated with these kinds of heuristics, which is the intuition behind Condorcet’s Jury Theorem and related studies. Miller brought this to the attention of the expert-identification debate in a simplified manner: (see Dietrich and Spiekermann 2023, chaps 2, 6.2; Miller 2013, 1304–5; Stegenga 2009, 654–60)

(...) sufficiently large groups of individuals in which there is a sufficiently large subgroup of individuals who have a higher than 0.5 probability to form a correct belief on a given matter will reach the correct decision on that matter by the method of majority voting. The application of these mathematical theorems to concrete real-world cases, however, is far from trivial. (Miller 2013, 1304 footnote 14)

Laypeople updating their beliefs in a Bayesian manner within a pool of putative expert testimonies should eventually reach the ‘correct’ belief if most testifiers are real experts and their expert testimony is more likely to be true than chance, which also fits my conceptualisation of expertise in Chapters 1 and 2.2. (see Dietrich and Spiekermann 2023, chaps 2, 6.2; Miller 2013, 1304–5)

Robustness. The state in which a hypothesis is supported by evidence from multiple techniques with independent background assumptions. (...) Robustness is often presented as an epistemic virtue that ensures objectivity. Champions of robustness claim that robust evidence can demarcate artifacts from real entities, counter the ‘experimenter’s regress,’ ensure appropriate data selection, and resolve evidential discordance. Consider the worry about artifacts: If a new technique shows x (an entity, process, relation, etc.), this might be due to a quirky aspect of the technique rather than a real description of x. *Response*: If x is observed with multiple methods, it is extremely unlikely that x is an artifact (Hacking 1983). (Stegenga 2009, 651 italics in original)

So, robustness is the idea that evidence or propositions derived from *multiple different* methods or sources are less likely to fail the same way, i.e. they are more likely to be driven by an underlying ‘real’ world. Standardly, this has been seen as a feature of statistical analysis, models, or experimental evidence, though in my heuristic context, it is the alignment of expert testimony (and indirectly their expert evidence) that increases the robustness of a proposition. This robustness also admits degrees, though Stegenga notes the difficulty of specifying these degrees precisely. However, to count as *distinct* techniques or pathways, they need to be properly independent of each other. Stegenga contends that this individuation between methods can be difficult in practice, yet vague distinctions would render robustness-style arguments less conclusive. (see Stegenga 2009, 651–56)

Moreover, if evidence from different techniques or sources displays varying degrees of consistency, congruity, or intensity, it is *discordant*, which means that such evidence is not or only partially comparable.³³ This is a problem since often: “We lack systematic methods for assessing and combining evidence, and without such methods, robustness is limited to a qualitative or intuitive notion” (Stegenga 2009, 652). Hence, there exists a tension between the *diversity* and the *concordance* of evidential pathways. The more alike methods and their interpretations are, the less they tend to be independent of each other and vice versa. (see Evans 2014, 238–39; Stegenga 2009, 654–60)

For my purpose, it is important to clarify whether or how discordant *testimonial* evidence is. I would characterise the latter as associated with two layers: One is the testimony, proper, in the style of x says y or not y. The other layer is the expert method each testifier draws from. The latter would, in the context of the expert-identification question, be very unspecified, allowing many kinds of methods, background assumptions, etc. After all, laypeople would not

33 Consistency demands methodological meta-standards that determine the analogous meaning of data in each compared method. Congruity refers to the basic coherence of the different scientific languages, requiring comparable background assumptions. Intensity refers to the statistical scope of a method, ergo what statistical frequencies pertain to epistemically in each method. (see Stegenga 2009, 654–56)

be able to filter out testimonies from discordant expert methods as they lack the required specialist knowledge. I think, however, we should see the discordance of original expert evidence itself as already intrinsic to the quality of evidence available, since our fallible heuristics cannot, by principle, be better than the best available knowledge. The interesting kind of robustness is the one *between testimonies*. Concordance here refers to the comparability of testifiers, i.e. being relevant for the question, and independence refers to experts testifying independently of each other. Testimonies from different (sub-) disciplines might be less than optimally concordant, as they might involve different norms of confidence and acceptable uncertainty. (see Stegenga 2009)

I think it now becomes clear how Goldman's and Miller's accounts rely on the rationale of robustness. Interestingly, Millers focuses on the criteria of concordance with his social calibration and apparent consilience of evidence conditions and only considers the independence aspect in recommending agreement with socially diverse testifiers. On the other hand, Goldman skips concordance by making it a matter of preselection by other heuristics and mostly concerns himself with the idea of belief-independence. By combining these analyses with Stegenga's study, I think to have established a more complete picture of criteria relevant to robustness and the difficult relationship between concordance and independence criteria. I also want to note again that robustness is a gradual property of testimonies and increases laypeople's reason to trust without requiring a strong consensus. (see A. I. Goldman 2001, 97–104; Miller 2013, 1300–1314; Stegenga 2009)

Goldman's belief-independence thesis, however, met critique by David Coady, Lackey and Miller, who claim that there are ways how dependent expert testimony can increase a layperson's justification to trust a position. Firstly, Miller states that the belief-independence thesis counters the common intuition that critical deliberation between individuals improves the outcome and that large groups can do research more effectively than small ones. The worry could be that by this principle some individual bogus theorists would be considered more epistemically contributing than professional collectives. Both David Coady and Lackey, furthermore, argue that also non-discriminatory reflection can justifiably increase belief if it results from a *meta-expert assessment*. Here, the additional testifier copies the testimony of the original expert *because* she knows about the latter's high competence. In this sense, dependent testifiers can function as an epistemic filter which preferably endorses trustworthy sources, if they have some meta-expertise or other relevant insight into the competence and honesty of the original testifier. Hence, if laypeople are confident that the additional testifier can function as a

filter, they can use this indirect assessment as a reason to increase trust in the original testimony. Note, that this meta-assessment does not need to have the same weight as a fully independent agreement from another expert. (see D. Coady 2006, 68–72; Miller 2013, 1304–5; Lackey 2013, 249–53)

Next to meta-assessments, Lackey also proposed that the additional testifier's non-discriminatory reflection can represent a *vouching for the truth* of the original claim. By this action, the additional testifier assumes a part of the reputational risk of the original testifier, since if the testimony turns out to be wrong, she might be held accountable as well. Lackey's argumentation, here, closely resembles Goodwin's analysis of how trust in experts is centrally mediated by a reputation economy. (see Goodwin 2011, 293–95; Lackey 2013, 249–53)

I want to argue that David Coady's and Lackey's arguments are not in contradiction to Goldman's belief-independence thesis but rather extend his analysis from being solely concerned with independent evidence on a proposition to the possibility of additional testifiers having independent evidence on the competence and honesty of experts, which in turn have independent evidence on a proposition. In refuting the belief-independence thesis, I think, David Coady, Lackey and Miller pay too little attention to the exact wording of Goldman and confound it with a rejection of anything but independent evidence. Goldman explicitly referred to 'non-discriminatory' reflection and the copying of testimony without assessing the truth of the original belief. Yet, when a layperson is justified to believe that an expert has an honest character and is competent in a certain domain her assessment of the expert is very well discriminating! It discriminates indirectly whether the expert can know about the issue and is inclined to tell the truth. Goldman even, explicitly considers how the judgment of meta-experts can rationally increase a layperson's trust in a putative expert. (see D. Coady 2006, 68–72; A. I. Goldman 2001, 97–104; Lackey 2013, 249–53; Miller 2013, 1304–5)

David Coady and Lackey miss the point, that the belief-independence thesis is simply about not counting the same evidence twice. This is underlined by Goldman's choice of extreme examples like that of a guru whose followers mindlessly echo his teachings. In more subtle scenarios, observing partial dependence can rationally motivate laypeople to give a reduced amount of weight to testifiers who report partly the same evidence. This rationale also applies to indirect assessments of additional testifiers. Someone mindlessly copying the critical assessment of someone else about the original testifier should not move the layperson to increase her belief in the trustworthiness of the original testifier. Also, Miller's worry regarding undervaluing research collectives is not serious, as neither is independence reduced to 0 in such

scenarios nor is it the only desideratum. People working together can follow a far wider scope in their projects. Nevertheless, David Coady's and Lackey's contribution is important, since they explicitly connect the relevance of indirect heuristic judgments to the heuristic of counting expert agreement, while Goldman solely listed access to independent eyewitnesses, experiments, and methods of reasoning as adding epistemic weight. (see D. Coady 2006, 68–72; A. I. Goldman 2001, 97–104; Miller 2013, 1304–5; Lackey 2013, 249–53)

What remains an open question is how much epistemic weight should be given to indirect assessments in the calculation. Apart from the general ambiguity plaguing the notion of 'increasing or decreasing belief' in the expert-identification debate, I would suggest that the epistemic value of indirect assessments is marginally decreasing. Certainly, Goldman's framework by nature already assigns diminishing returns due to its Bayesianism, but the decline of the additional value of indirect assessments seems to be even steeper. After all, one cannot gain certainty in one's competence and honesty beyond deeming a person perfectly suitable and trustworthy. Indirect testimonies, then, are always secondary to direct evidence on propositions, rendering 30 direct plus 20 indirect testimonies preferable over 10 direct plus 100 indirect ones. (see C. A. J. Coady 1994; D. Coady 2006, 68–75; A. I. Goldman 2001, 97–104; Lackey 2013, 257–60)

3.3.3 Indicators for robust testimonial evidence

In the previous chapter, I discussed relevant criteria regarding which kinds of agreement increase the robustness and thus trustworthiness of expert testimony. What is missing still is how laypeople could learn about the fulfilment of these criteria. Indeed, how concordant and independent agreement is, is not self-evident. After all, Stegenga highlighted the difficulty of finding ways to sensibly compare different data and methodologies and demarcate distinct evidential pathways. While this is a matter of puzzling concern for (scientific) expert debates, the demands for laypeople need to be different. As I mentioned in the previous chapter, robustness in the expert-identification question concerns the testimonies of experts and not (directly) how they contrived them, which would be inaccessible to laypeople anyway. (see Stegenga 2009)

In this vein, a concordant expert agreement is indicated by anything that makes it more likely that a putative expert has the *relevant* kind of expertise and is addressing the question of interest. This could be a seemingly fitting training background like having finished an apprenticeship as

a car mechanic for testimonies in car repair or having studied medicine to be a medical expert. Moreover, other licenses granted by meta-experts like certificates or continued education for relevant specialisms of, for instance, law or counselling can signify this. Even work experience can entail the relevant kind of alignment since concordance means comparability and, in itself, not competence or trustworthiness. If putative experts share such indicators, they are comparable, without needing to share identical indicators. Again, on a testimonial level, we compare people who say proposition x is a good or bad idea not the concordance of their data and methods. (see Stegenga 2009)

Independent testimony, on the other hand, is indicated by anything that makes it more likely that an expert has reached their opinion by an independent reasoning process. Here, the argumentation regarding Goldman's belief-independence thesis fits the level of testimonial lay comparison well. Independence, then, is fallibly indicated by experts not working closely together in the same organisation, department, team, or on the same project. Moreover, different sub-disciplinary backgrounds, like oncology and immunology in medicine, and theoretical commitments like proponents of Neoclassic or Keynesian economics signify that the agreement is a product of the convergence of *different* perspectives. (see A. I. Goldman 2001, 97–104; Stegenga 2009)

Finally, Miller's third 'indicator' of socially diverse agreement is a real indicator as such information is easily identifiable by laypeople who are normally socially encultured in a society. Such could be gender, class, nationality, etc. Our social positioning including our individual experiences commits us to a certain perspective on the world while at the same time excluding other possible viewpoints. As already mentioned in Chapter 3.1, social identities can form biases that constrain our actions and thinking. Millers' idea is that if several putative experts of diverse social identities agree on a proposition, then it is less likely that this consensus is affected by idiosyncratic biases of one social identity.³⁴ (see Collins and Weinel 2011; A. I. Goldman 2001, 104–5; Miller 2013, 1311–14)

In his paper, Stegenga worried that the independence of evidential sources increases discordance, and concordance increases their dependence. After all, having been trained in the same institutions and methods makes it also more likely that experts share the same background

34 Note that this is a *ceteris-paribus* concern and social diversity indicates trustworthiness by robustness (against biases) and not competence. The sociological literature has also seen a debate on whether a diverse selection of experts epistemically outperforms a homogeneous group, including various definitions of diversity. This is a different matter and not the concern of my thesis. (see Hong and Page 2004)

assumptions. As Evans notes, this makes it more likely that they are pre-disposed to agree on framing a question in a certain way that will not be critically reflected. He calls this fixation on a frame the opportunity cost of becoming a specialist. I agree that this tension exists but does not need to defeat robustness. Looking closer at how concordance and independence are indicated on a testimonial level, these criteria do not cancel each other out but extend each other. Within a space of relevant testifiers different theoretic or methodological commitments increase independence more than they create discordance as testimonies rather simply add to testimonies. Vice versa, considerations of relevant expert testimony regarding a specific domain can be done more or less exclusively. On one side, too much tolerance regarding different perspectives entails the risk of including irrelevant, low-quality contributions diluting the epistemic value of the heuristic. On the other hand, being too exclusive might reduce the potential for the convergence of different backgrounds thus reducing the independence and robustness of testimony. Without framing robustness in the same testimonial manner as myself, Evans (and Collins) call this ‘the problem of extension’. Accordingly, robustness is highest at the optimum between including too few and too many kinds of testifiers. This plus the diligent use of direct heuristic judgments constitutes the relevant selection and weighting of concurring testimony to compare the agreement towards a proposition. From this point on more relative agreement to a position gradually makes this testimony more robust and, *ceteris paribus*, more trustworthy. This is fallible but should be seen as part of a heuristic framework that navigates the best possible decisions within a difficult epistemic environment for laypeople. (see Collins and Evans 2009, 113–15; Evans 2014, 238–45; Stegenga 2009)

A remaining problem is that laypeople have limited time and energy resources and, thus, can assess only a limited sample of testimonies. Nevertheless, even comparing some testimonies makes reaching the correct belief more likely and thus is a sensible heuristic. One might still worry, what if the information up until the limit of laypeople’s ability to process testimonies is not representative of the entire available data? Repeating a mantra of my thesis, this is, unfortunately, a possible scenario, approximatively however, a limited amount of research is more likely to lead to a better-justified position than none, especially if people stay somewhat self-critical regarding their ignorance. Further, Collins and Weinel emphasise, that the state of the art of (scientific) expertise changes over time. Accordingly, what was once open to debate might be settled later (until new evidence emerges). This shows that laypeople need to be aware of whether the expert testimony interacts with the *current* debate. (see Collins and Weinel 2011, 408–11; Dietrich and Spiekermann 2023, chaps 2, 6.2; Miller 2013, 1304–5)

To summarise the last two chapters, ‘counting agreement’ is indeed a valuable indirect heuristic judgment that justifies laypeople to increase their trust in a position. The more agreement a proposition has relative to others the more robust is it for laypeople, *ceteris paribus*. More precisely, the epistemic weight of each agreeing testifier depends on its independence and concordance. Though the properties of independence and concordance often are in tension with each other, there might be an optimum between both poles that assures the most epistemic weight. The agreement of testifiers are social facts and thus easily accessible indicators. Concordance is indicated by experts seemingly relevant to the question of interest, i.e. having a suitable education, licenses or experience, whereas independence is indicated by different methodological and theoretical commitments and the diversity of their social identity backgrounds. Finally, robustness is at its optimum where the extension of the compared field is properly balanced, though the associated indicators remain fallible.

3.3.4 Strategically utilizing multiple testifiers to disincentivise dishonesty

In the previous section, I discussed a heuristic of passively counting and evaluating agreement between putative experts, but this need not be the only way to take advantage of testimonies from multiple experts. Similarly to the assessment of incentives in Chapter 3.1.3, laypeople could engage strategically with putative experts. Using economic game theory with multiple testifiers might give insights into whether this epistemic environment is more beneficial for laypeople. Accordingly, I am considering the extended case of one receiver being able to observe messages from multiple senders (i.e., at least two), which changes the strategic position of the former. This strand of modifications of the original cheap talk model is growing rapidly as scholars add new arguments, modelling assumptions and further complexities. I choose to focus on one of the more central and stable insights of this debate, which could more reliably inform my applied research question without losing myself in detail. (see Ambrus and Lu 2014; Crawford and Sobel 1982; Sobel 2013, 317–21)

In this section I want to specifically analyse the following strategy in unidimensional state spaces (i.e. unidimensional propositions and testimonies), which is prominent in the many-senders literature:

Having multiple Senders allows the Receiver to check facts. If the Senders have access to identical information, the Receiver expects identical messages. If the set of actions is sufficiently rich, then the Receiver can respond to conflicting messages with an action that punishes both Senders, inducing informative reporting in equilibrium. (Sobel 2013, 317)

In their 2014 paper, Ambrus and Lu want to provide formal support for this intuition by extending the standard single-sender model of Crawford and Sobel with an additional sender and checking whether a fully revealing Bayesian Nash Equilibrium is possible. Now, two biased senders each learn individually about the true state and send a message to the receiver, who, then, observes two messages. Also, *prima facie*, the senders would have each an interest in messaging their individual ideal testimonies and abuse the information asymmetry. Since the senders' bias is common knowledge, the receiver would be able to deduct that, which in turn can be anticipated by the senders and is not in their interest. Conditional on, again, not having too large biases, the senders can partition their type space into several intervals, which are called cells here. In this many senders case, Ambrus and Lu propose that the receiver can detect whether testimonies conflict, i.e. are not identical, and infers that at least one sender must have deviated from the true state. If the bias is small enough and the action space big enough, the authors assume that the receiver can choose an extreme action if she observes conflicting testimony. The only way for each sender to avoid this extreme outcome is to message the true state of the world, as any non-coordination with the other sender would result in an outcome worse for them than truth-telling. Hence, Ambrus and Lu show that within this model design, fully revealing equilibria are possible.³⁵ (see Ambrus and Lu 2014, 174–81; Sobel 2013, 317–21)

The outcome of this modelling endeavour contrasts the results of the original single sender cases considerably, where fully revealing equilibria were impossible when the sender had any bias. However, it also seems to contain unrealistic assumptions, due to following a weak notion of perfect Bayesian Nash equilibria. The receiver's strategy of choosing an extreme action if the senders' messages conflict, might be an incredible threat since, off-the- equilibrium-path, the receiver would be better off by optimizing according to the imprecise message the senders would otherwise communicate than bearing the consequences of her own extreme action. It is questionable, then, whether this can be a rational prediction for strategic interaction. Moreover, anticipating the strategies of multiple senders and their bias might demand a lot of cognition

35 Sobel outlines Ambrus and Lu's (2014) model in the following manner: "Break the interval of types into n^2 intervals. Label the intervals (i, j) , where $i = 0, \dots, n - 1$, and $j = 1, \dots, n$. Sender 1 reports $m_1 = j$ when she observes θ in interval (i, j) . Sender 2 reports $m_2 = j + i \bmod n$ when she observes θ in interval (i, j) . The Receiver responds optimally to the message." (2013, 318–19)

from laypeople, which puts additional stress on the rationality assumptions in this model. (Ambrus and Lu 2014, 174–78; Crawford and Sobel 1982; Sobel 2013, 317–21)

Despite this, Ambrus and Lu's model shows a way how we could think abstractly about the presence of multiple experts could improve the epistemic position of a layperson. So far, I deem the assumption that laypeople could detect experts' disagreement as not controversial, since, under provisional assumptions, experts are perfectly competent. Disagreement, therefore, could only be the result of non-epistemic, detrimental motives of at least one expert. (see Ambrus and Lu 2014; Sobel 2013, 317–21)

This strategy demands a large action space, which translates into bargaining power. The more restricted it is the fewer extreme options are available to potentially punish the experts with. An example of such a punishment could be the avoidance of misbehaving experts by choosing a competitor or the warning of other laypeople to incur an income loss. Using Almassi's terms, greater (rationalizable) abilities of laypeople to punish misbehaving experts reduces the latter's 'opportunities' to bias advice. In real-world contexts, however, I doubt that such individual manoeuvres can be sufficiently extreme since experts often have enough other clients, sources of funding, or reputation to render them inconsequential. Perhaps provoking a massive lawsuit could suffice. After all, heuristics like these depend on strategic bargaining power which in turn depends on contextual circumstances like individual capital, laws, control institutions, traditions etc. Moreover, typical laypeople are constrained in time, money and education, which further raises doubts about their ability to sufficiently threaten putative experts or anticipate their complex strategic behaviour. Furthermore, if the experts' and laypeople's interests are misaligned too much, laypeople will not be able to find a point extreme enough to threaten the experts since the latter's ideal point is already substantially deviant. Again, there is a limit to how biased experts can be, for laypeople to be able to receive informative testimonies. (see Almassi 2017, 207–14; Ambrus and Lu 2014; Sobel 2013, 317–21)

There might be another situation where laypeople could not induce a fully revealing equilibrium independent of the problem of incredible threats. Goldman argues that when the entire expert community shares the same bias, avoiding its detrimental effects becomes difficult. Experts might, for example, have an incentive to collectively protect the reputation of their discipline, as this would affect them all the same. Boaz Miller raises the example of experts manufacturing a consensus to mask disagreement thereby increasing the appearance of certainty in their discipline to feign higher precision and reputation. When experts are perfectly informed and share the same bias, they might prefer to send two identically biased messages. In this case,

the senders' messages are biased without conflicting, rendering the testimonies inscrutable by the receiver within this heuristic. At best, then, partially revealing equilibria are possible. (see Ambrus and Lu 2014; Miller 2013, 1299; A. I. Goldman 2001, 105–9)]

4. An interdisciplinary apprehension of personal-finance expertise

In the first part of my thesis, I have critically assessed the general debate on how laypeople can use heuristics to rationally trust experts. In the process, I focused on what I consider the most promising heuristics and complemented these with an economic game theory perspective, which filled gaps in the agential possibilities the philosophical debate had overseen. Now beginning with the second part of this paper, I will try to apply the insights of the general debate to the specific case of personal-finance expertise. Firstly, this concretisation helps in critically reflecting on our previous discussion as it prevents us from protecting arguments with vague abstraction and therefore enables us to better gauge the epistemic effectiveness and constraints of heuristics.

Secondly, personal finance is of intimate relevance for citizen as it concerns decisions regarding their individual (monetary) funds and thus most directly affects how they can fulfil their needs and desires. Certainly, decisions on climate policies in democratic contexts should be similarly important to us as their consequences too will deeply affect our lives. Nevertheless, personal finance decisions, but perhaps also dating or fashion advice, might seem closer to heart since one's own decisions have a higher impact on one's own well-being. Also, an individual, *prima facie*, has proportionally more control over the outcomes in these matters, hence, more depends on her own level of knowledge. This discretionary intimacy motivates people to consider these questions and seek experts to improve their epistemic situation in this regard. Therefore, knowing how to rationally trust personal-finance experts might be of special interest to laypeople.

Moreover, the idea of more deliberately managing one's own money by engaging with financial markets or taking out personal loans has risen in the last decades. Starting around the 80s, especially in the United States, Great Britain and Australia, personal finance topics have received continuously more attention. Roei Davidson connects these developments with a risk shift from collective-institutional actions to individual responsibility, which accompanied a transition from a corporate-liberal to a neoliberal economic order. In studying the emergence of personal-finance magazines, Davidson describes the genesis of a social identity as a self-

reliant and resourceful citizen who is responsible for her own financial fate. While the motivation to proactively engage with these topics is rising, Inderst and Ottaviani cite empirical data from the mid to late 2000s that show a high dependence on financial advice from classical institutions like house banks, which are often subject to corrupting incentives. Albeit the financial literacy of the public seems to be rather humble, trends show a migration from classical institutions to more self-involving approaches mediated by Fintech tools. These developments make it especially interesting and necessary to think about how laypeople could heuristically navigate personal-finance expertise. (see Beek 2022; Davidson 2012; de Bruin et al. 2023, chap. 4.2.1-4.2.2, 5.1; Inderst and Ottaviani 2012, 393–96)

Thirdly and most importantly, this analysis can show us how far applicable the expert-identification debate is to a social-science-informed, non-scientific cognitive expertise like personal finance, as I will demonstrate in the following chapters. Presently, I want to argue that the philosophical expert-identification debate has, so far, not so much discussed the phenomenon of cognitive expertise in its full range but noticeably focused on (natural) scientific expert testimony. For one, this is suggested by the major use of natural scientific cases as examples and case studies in contribution that centrally respond to the debate. Consider, for instance, the papers of, Grundmann, Almassi, Lane, Miller, Collins and Weinel who primarily referred to issues of climate sciences, environmental pollution, physics, and medicine. (see Almassi 2017, 210–11; Anderson 2011, 149–53; Collins and Weinel 2011, 409–11; Grundmann 2024, 17–18; Lane 2014; Miller 2013)

Occasionally, contributions have addressed *singular* aspects of the expert identification in contexts of the social sciences, especially economics, like C.A.J Coady, Martini, and Evans. Foremost still Marcel Boumans and Carlo Martini investigated potential differences and peculiarities in their anthology on *Experts and Consensus in Social Science*. (see C. A. J. Coady 1994, 277–304; Boumans and Martini 2014; Evans 2014; Martini 2014a, 12; 2014b; Martini and Boumans 2014)

Next, some authors did not exclude non-scientific expertises but also remained so general to pre-empt the occurrence of any possible differences like Duijf, Goodwin and David Coady. (see D. Coady 2012; 2006; Duijf 2021; Goodwin 2011) Personal-finance expertise, for instance, has been used as an example but has not been studied thoroughly beyond that. (e.g. Collins 1992, 167; Collins and Weinel 2011, 402–3; Duijf 2021, 9293; A. I. Goldman 2021, 91; Martini 2020, 117). Such a generalisation would be justified if the discussed heuristics were equally applicable and effective in every kind of expertise.

Given this prioritization in the philosophical expert-identification debate, it is not surprising that factors like intrinsic uncertainty, non-scientific expert control institutions, high ambiguity or localisation have received scant attention. Moreover, it might explain the optimism in the ability to avoid biases and lack of nuance accordingly. What is missing is an exhausting paper on the putative differences in the lay assessment of testimony between different kinds of expertise. My thesis should be a contribution to the case of personal-finance expertise, where I discuss how applicable heuristics are in this context. The first step, however, is to describe what personal-finance expertise is and what it entails. (see Lane 2014, 105–11)

4.1 The discipline of personal finance

As the name suggests, contrary to finance which concerns itself with the money management of nation-states, public bodies, private firms or banks, personal finance revolves around the budget of an individual, family, or household. In its modern form and under the banner of ‘personal finance’, this discipline is quite young: “Prior to 1990, the majority of researchers interested in the topic of personal finance considered themselves to be family economists, consumer economists, consumption economists, household resource management specialist, or consumer educators.” (2007, 62), as Jane Schuchardt et al. review the history of the discipline. (see Grable and Chatterjee 2022b, 2–10; Schuchardt et al. 2007, 62–66)

Therefore, most scholars contributing to personal-finance topics and journals still come from an academic background in economics, finance or family and consumer sciences. For the longest time, this field of research has not received attention from mainstream economics, business, or finance, until John Campbell wrote a macroeconomic paper on household finance. In the two decades following this increased interest, the different disciplines and research programs studying these questions began to form this interdisciplinary field under the overarching domain of ‘personal finance’. (see Campbell 2006; Grable and Chatterjee 2022b, 3–10; Schuchardt et al. 2007, 62–66)

Drawing on previous work from Jane Schuchardt et al. and John Grable and Swarn Chatterjee propose the following definition of personal finance:

Personal finance is the study and application of concepts, tools, and techniques associated with the planning and management of personal and household financial activities, including generating income, managing spending and debt, saving, investing, and protecting sources of income and assets in a way that informs practices and policies designed to enhance the wellbeing of individuals, families, and households. (Grable and Chatterjee 2022b, 10)

The service function of discipline is, accordingly, to improve the financial capabilities, literacy and security of individuals or households, which is also summarised as ‘financial well-being’ in this expert community. (see Grable and Chatterjee 2022b, 7–10; Schuchardt et al. 2007, 63–65) To increase the financial well-being of their clients, personal-finance experts engage in:

“(…) the study and use of tools such as financial statements, checking and savings accounts, investment products and strategies, and daily financial management products and services (e.g., management of debt, homeownership, bankruptcy, risk-management techniques, tax planning, retirement planning, special needs planning, and estate planning). (Grable and Chatterjee 2022b, 7)

So far, I have characterised the domain of personal finance in general terms, but it can be further divided into a scientific, a pedagogic³⁶ and a practical service expertise. It is the latter, which is of central interest to laypeople as it constitutes the profession that directly engages with laypeople and advises them on personal finance issues. These are the kind of experts whose lay assessment I want to discuss and what I am, from now on, addressing when I speak about personal-finance experts. (see Grable and Chatterjee 2022b, 7–10)

Personal finance is also a cognitive expertise. Questions like ‘how I should invest my savings to receive the most returns’ or ‘what the most efficient way is to reduce my debts’ depend on market structures along with behavioural attitudes and are propositional in nature. The expert is consulted because she knows about the complex market behaviour relevant to individuals, their legal context, which strategies are available, and which serve typical clients’ interests best. Only experts with real competencies could fulfil this function and are of relevance to laypeople, as argued in Chapter 1. While these personal-finance experts are not scientists themselves, they draw on relevant scientific knowledge as discussed earlier. This renders them knowledge intermediaries who apply a discipline’s scientific body of knowledge to specific strategies that suit laypeople’s interests and abilities. In conclusion, personal-finance expertise is a normative

36 In their public service function this profession might have some overlap with educators who disseminate basic personal finance literacy. To this degree, I think we could see these endeavours as equivalent but not beyond that, since, as discussed in Chapter 2.2, laypeople typically do not have the resources to learn more than the utmost basics.

and cognitive expertise, where advice is being transmitted to the receiver.³⁷ (see A. I. Goldman 2001, 91; 2018, 4; Grable and Chatterjee 2022b, 7–10; Schuchardt et al. 2007, 63–65)

Often personal-finance experts further specialise themselves in a certain range of topics or tasks, but since *personal* finance is already a rather specific area, the additional differentiations are less extensive. Typically, a personal-finance expert ought to guide an individual with a range of tasks, that more holistically suit the needs of the client. Here, rather common ways of specialisation concern different financial needs of laypeople and their legal context. The work of an expert regarding a highly indebted client differs substantially from that of a wealthy heir who needs advice on investing her fortune. (see Grable and Chatterjee 2022b; Schuchardt et al. 2007)

In this regard, I want to distinguish the specific practice of (personal) finance forecasting from other more ‘basic’ expert practices in this discipline. The latter are more heterogeneous assortments of tasks like helping to structure debt, managing money and income daily, finding proper ways to save money, planning for future expenses, avoiding taxes, choosing suitable insurance, etc. These I will call ‘basic personal-finance expertise’ in this thesis. From these expert practices, I want to single out those that depend on financial forecasting. Here, I am referring to expert’s attempts to predict the future development of financial parameters like the performance of market indices, currencies or individual financial products like stocks or other securities. Knowing about their future behaviour would be extremely relevant for the client’s financial planning and hence well-being. Despite its ubiquitousness in the profession, serious doubts regarding its effectiveness exist. This is the reason, why I want to analyse this part of personal-finance expertise separately in Chapter 4.3. (see de Bruin et al. 2023, chap. 4.2.2; Grable and Chatterjee 2022b, 7–10; Tetlock 2017; Schuchardt et al. 2007, 62–82)

Another relevant distinction for the expert-identification question is the medium of testimony transmission, which I call personal or impersonal. Firstly, laypeople can receive expert advice impersonally by reading books or magazines, watching tutorials, or listening to podcasts. Here, the putative expert might have a target audience in mind and the layperson might filter for contributions relevant to her situation, yet such advice tends to remain less personalized to laypeople’s needs. After all, experts cannot address all possible circumstances

³⁷ The literature on expertise in social sciences also sometimes conceptually detaches ‘expert-judgment’ from the use of (formal) decision-making methods by a scientific expert, whereas the former only features a subjective assessment based on an epistemically privileged position. I do not make this distinction in this paper and consider expert testimony simply as the result of any means available to the expert. (see Reiss 2014, 143–50; Tetlock 2017; Watson 2020, 41–48)

at the same time. Furthermore, a layperson might lack some tacit understanding of what could be important to her situation and would consequently not look for that. On the plus side, more distance to the putative experts might make it easier to compare them and their advice, since books or videos tend to have lower entry costs regarding time and money. Also, emotional attachments, social pressure, or sunk-cost fallacies are less likely to affect laypeople. Note, that I am speaking about rough tendencies here as the market for personal-finance advice comes in many forms. (see Davidson 2012; Beek 2022; Grable and Chatterjee 2022b)

Alternatively, testimony can be transmitted personally, which often involves eye-to-eye conversations, which are structured by contractual obligations. Laws guide which kind of advice is suitable enough for clients according to their preferred risk behaviour, financial situation, goals, etc. Unsuitable advice or lack of transparency might have legal consequences for the adviser. Advisers, however, are often associated with a financial product provider and might have an interest in pressuring the laypeople into buying products that are more beneficial for this provider than for themselves. Personal settings, nevertheless, also allow the personal-finance expert to assess the situation of their client and advise them on their individual needs, while the laypeople have the chance to immediately clarify questions. The mode of payment here depends on whether the advice is independent or dependent. In the latter case, a putative expert is typically paid by commissions from product providers for recommending and selling their financial products. Independent advice, instead, is often paid by a fee upfront. Having the first part of my thesis in mind, we might already see that this is consequential for how laypeople will be able to trust experts. (see Inderst and Ottaviani 2012, 394–96)

Apropos legal protection, in the aftermath of the financial crisis of 2008 and the posterior European debts crisis, the EU imposed a new directive (MiFID-2³⁸) onto its member states. This framework is an example of how lawmakers intend to increase consumer protection and expert liabilities by regulating the market for lay consumer financial products within their jurisdiction. Such laws might determine the required training of advisers, the associated control institutions, the structure and transparency of counselling sessions, the design of products targeted towards laypeople, etc. Originally MiFID-2 envisaged to ban dependent advice, i.e. commission-taking, to protect clients from putatively corrupt incentives. Apart from the Netherlands and back then Great Britain, all other member countries of the EU opted out of this

³⁸ MiFID-2, i.e. Markets in Financial Instruments Directive 2, has been active since 2018 and replaced its predecessor MiFID-1 in regulating and harmonising the European market for financial products. (see Directive 2014/65/EU; also see <http://data.europa.eu/eli/dir/2014/65/oj>)

aspect. In Austria, for example, dependent advice is the most common variety. (see Beek 2022) (see Directive 2014/65/EU)

Finally, I want to address the localised character of personal-finance expertise.³⁹ Next to laws regulating the design and recommendation of financial products, personal-finance advice also depends on national tax law, the economic and fiscal policies of a jurisdiction and local financial customs. After all, these local characteristics likely shape which strategies are best for the layperson to fulfil her goals. This renders personal-finance expertise a partially localised expertise, i.e., a proper expert in a US context could not retain her epistemic authority if she were spontaneously moved to Serbia. Though many aspects might remain relevant, the migrated expert would need to learn the local context to be sufficiently competent again. The general rationales behind the heuristics to assess their trustworthiness, however, which I am discussing in this thesis, remain applicable wherever there are personal-finance experts of the kind I described in this chapter. The outcome of the use of these heuristics might differ on the other hand. If the legal protection in a jurisdiction is only weakly developed, laypeople, *ceteris paribus*, have fewer reasons to trust experts. Similarly, if meta-expert control institutions are missing, laypeople might lack the means to detect and punish biased testimonies and need to focus more on other heuristics. (see Grable and Chatterjee 2022b; Schuchardt et al. 2007)

4.2 Personal finance as a social science

This interdisciplinary domain is mainly informed by economics, psychology, and sociology, which firmly embeds it into a social science context. Understanding to what degree scholars can be epistemic authorities in this discipline is consequential for what laypeople should expect from practical experts, who draw on this specific body of knowledge. In the next chapter, I will more precisely discuss the epistemic scope of different aspects of personal-finance expertise, especially economic forecasting, but first I want to address a few general considerations regarding the scientific character of the field. (see Schuchardt et al. 2007, 63–65)

Social sciences like economics study social phenomena and their interrelations, involving the search for regularities and patterns in human interactions. This differentiates them from natural sciences, who study the impersonal mechanisms of nature instead. While scientists of the latter seek causal explanations, the research object of social sciences adds another

39 Do not confuse the discussion here with Quast's 'localisation problem'. (see Quast 2018, 408–11)

dimension: Social agents can have ‘reasons’ for acting. Daniel Hausman argues that explanations in terms of reasons are substantially different from purely causal laws, since reasons can justify actions besides merely explaining them but must be intelligible to the agent who invokes them. Drawing from Donald Davidson, Daniel Hausman asserts that for reasons to explain an action means to cause it. Ergo, social sciences like economics, which studies the behaviour of participants in a market, require an account of agential rationality, which itself is subject to debate in economic methodology. Moreover, reasons that causally move agents can be the result of mistakes and ignorance, this expands possible causes to counterfactual beliefs. (see Hausman 2021, 2.2)

The concept of rationality introduces a normative element to the study of economic behaviour. It is not descriptive of how agents typically act but it proposes minimal assumptions of what it means to ‘properly’ act, as opposed to mistakenly or unreasonably. For example, that agents rank expected outcomes of actions according to the utility or benefit they bring them, that this ranking is complete and transitive and that they prefer the feasible action that entails the most utility. Such assumptions are often used as fixed starting points to model more complex interactions between agents under the presupposition that ‘in general’ and ‘ceteris paribus’ ‘real’ humans approximate such rationality well enough so that these models can predict or explain tendencies in ‘real’ human interactions. Investigating the reasons and rationality of human behaviour, Hausman argues referring to Max Weber, does not just amount to predicting behaviour but also provides understanding from the internal standpoint of the agent. In Hausman’s words: “This (and the closely related recognition that explanations cite reasons rather than just causes) seems to introduce an element of subjectivity into the social sciences that is not found in the natural sciences.” (Hausman 2021, chap. 2.3; Weber 2011, chap. 2) (see Hausman 2021, chap. 2.2-2.3)

This element of subjectivity is of central interest to Julian Reiss, who considers the question of objectivity to be at the heart of methodological debates in economics. While also natural scientists need to ‘objectively’ justify their work by differentiating the essential mechanisms of phenomena from incidental elements and need to render their scientific method reproducible independent of personal elements, social sciences have an issue in objectifying explanations of social interactions, due to the intrinsic normativity, that is unique to them. Drawing on Weber again, Reiss contends that there is no a-perspectival social science and one’s social position always enters the analysis. (see Reiss 2014, 139–40; Weber 2011, chap. 2)

First, reality is too complex to admit of full description and explanation. So we have to select. But, perhaps in contraposition to the natural sciences, we cannot just select those aspects of the phenomena that fall under laws and treat everything else as “unintegrated residues” (...). This is because, second, in the social sciences we want to understand social phenomena in their individuality, that is, in their unique configurations that have significance for us. (Reiss 2014, 139)

Since social scientists need to make normative decisions regarding which factors are relevant, which reasons explain behaviour, what is rational and what is mistaken, this potentially allows multiple reasonable ways to analyse and frame patterns in social interactions. Analogously to the value-free ideal in science, social scientists might ‘objectify’ this aspect by postulating effective strategies to reach a goal without evaluating the desirability of that goal. (see Reiss 2014, 139–40)

This approach of restricting the influence of values to those instrumentally necessary is not easy, since “economic theories bear so centrally on people’s interests, there are bound to be ideological biases at work in the discipline” (Hausman 2021, chap. 2.1). Compared to basic research in physics or geology, insights from social sciences like economics are immediately consequential for the scientists themselves and many other stakeholders. More contrastingly, testimonies from the social sciences may be beneficial to some partial interests even or especially when they are wrong. This is a factor extending the set of potentially detrimental incentives. (see Hausman 2021, chap. 2.1; Reiss 2014, 139–40, 148–49)

As already mentioned, in social sciences it seems to be especially difficult to sieve out those events and relations that cause or explain a phenomenon from those that are incidental or artefacts of research. It is a common theme in the philosophy and methodology of economics to attribute a vexingly high complexity and elusiveness to economic phenomena that impede scientific progress. Reiss, for example, mentions the issue that economic factors often do not combine additively but interactively, which renders isolating economic laws difficult. Addressing this issue, Hausman conjectures social phenomena might be too irregular to be explained by a framework of laws and theories. High complexity notwithstanding, human interactions also seem to feature regularities, which might prevent us from endorsing scepticism prematurely. Even if economists’ use of idealisations, toy models and *ceteris-paribus* clauses are a matter of continuing debate, their purpose is to navigate this high complexity. (see Boumans and Martini 2014, 1–4; Hausman 2021, chap. 2.3; Reiss 2014, 140–41)

In this vein, Boumans and Martini cite Alfred Marshall, who emphasises that since economic mechanisms enfold within a messy causal background, the tendencies that can be captured by economic analysis feature a heightened inexactness and fallibility:

The laws of economics are to be compared with the laws of the tides, rather than with the simple and exact law of gravitation. For the actions of men are so various and uncertain, that the best statement of tendencies, which we can make in a science of human conduct, must needs be inexact and faulty. (Marshall 2009, 26)

Following the discussion on uncertainty in Chapter 2.2, testimonies in economics seem to involve a higher degree of intrinsic uncertainty. At least when we compare the main body of insights from economics and physics respectively, since at the outskirts of research findings are always quite uncertain. (see Boumans and Martini 2014, 1–4; Lane 2014)

Whether the boundaries between natural and social sciences are sharp or even fundamental might be debatable. After all, some issues in natural sciences are also very intimately consequential for people, involve normative proposals and evaluations like medicine. Moreover, natural sciences know also multiple methods and non-social phenomena can be complex as well. Whether normativity and complexity in social sciences is something fundamentally different from other sciences cannot be determined in this paper. What matters for the expert-identification question in the context of personal finance is what its social scientific sources entail for its epistemic authority. (see C. A. J. Coady 1994, 287; Hausman 2021, chap. 2.3; Reiss 2014, 139–40)

Notwithstanding this issue of demarcation, Boumans and Martini raise our attention to a more overt difference between social and natural sciences which might be a consequence of aspects discussed before. Social sciences show a substantially higher amount of disagreement within their discipline about what constitutes legitimate scientific problems and methods than typical natural sciences. Ergo, social sciences have a far higher diversity of theories and methods.⁴⁰ (see Boumans and Martini 2014, 1–4)

40 Boumans and Martini also argued that this lack of unified assumptions and methods compared to natural sciences inspired Thomas Kuhn to think about paradigms in science as temporary states, where a unified understanding allows a phase of normal science before new incompatible evidence triggers a scientific revolution. Kuhn argues that the existence of such paradigms is an indicator of a mature science. The evaluation of social sciences to be immature or not ‘properly’ scientific at all is resembled by the common distinction of the hard or exact natural sciences versus the soft, inexact social sciences. Also, C.A.J. Coady refers to these sciences as nascent and argues that, while having developed theory and proper empirical constraints, they lack methodological confidence and stability of results in contrast to mature sciences like physics. (see Boumans and Martini 2014, 1–4; C. A. J. Coady 1994, 287–89; Kuhn 2012)

Social scientists use a variety of tools for their investigations. In fact, probably no scientist has exactly the same toolset as another scientist. They gather evidence from different sources, under different conditions and possibly with different instrumentation. They share at times different beliefs on the efficacy of a particular method or validity of a theory, even when they share a general paradigm. Particularly in the social sciences, scientists are both the locus where different sources of evidence aggregate, in no pre-determined way (...) and also a direct source of evidence that comes in as intuitions, background knowledge, knowledge of initial conditions, and so on. (Boumans and Martini 2014, 4)

Indeed, Reiss predicts that economics is not going to overcome this heterogeneity anytime soon as its potential for disagreement is a consequence of its social object of research: Entanglement of positivity and normativity allows reasonable disagreements about facts even when data and observations are being shared. The complexity and elusiveness of economic phenomena further promote the eligibility of different approaches. (see Reiss 2014, 149–50)

Interestingly though, out of all social sciences economics seems to, at least superficially, resemble natural sciences the most. Economics has a relatively unified mainstream of methods and background assumptions that are occasionally but rarely abandoned. Certainly, heterodox traditions of economics exist but they do so separately on the fringes of the discipline and are developed independently of the mainstream. Moreover, some basic laws and relations are stably accepted by most scientific experts. Hausman further emphasises that economic theories have been axiomatized and thereby closely resemble natural sciences. In Thomas Kuhn's sense economics indeed seems to approach maturity but Martini argues that such a verdict is premature. (see Hausman 2021, 2.3; Martini 2014b, 120–23; Kuhn 2012)

One argument that a unified body of knowledge exists in economics is that popular textbooks have mostly uniform contents. Martini, however, contends that economic textbooks only appear to be relatively homogenous because they exclude a good number of subjects like the history of economics, computational or agent-based modelling, or experimental economics. Martini further considers the argument that economics has developed a common mathematical language which is of paradigmatic quality. Nevertheless, he criticises that this mathematical language has not helped in delineating a field of scientific problems on which all methods converge or are especially successful. (Boumans and Martini 2014; Martini 2014b, 121–26; Reiss 2014; Stegenga 2009)

I think Martini is too quick to completely dismiss both the textbook and the language argument. Uniform textbooks at least suggest that there are some topics and ideas scholars in the discipline deemed as stable enough to teach to students. Moreover, even if a common

mathematical language is not expressive enough to constitute a paradigm, it can increase the concordance of research by increasing the congruity of different scientific contributions. Following Stegenga, this would improve the epistemic value of robustness analysis in this field. After all, Boumans, Martini and Reiss do see a lot of practical, epistemic benefit in comparing evidence from multiple methods and sources, as I have analysed in Chapter 3.3.2 and will apply in Chapter 5.3. (Boumans and Martini 2014; Martini 2014b, 121–26; Reiss 2014; Stegenga 2009)

Even though economics is the main scientific source of personal-finance expertise, we should not forget that it is also majorly influenced by insights from psychology and sociology. The latter sciences might, for example, inform the expert more deeply about lay attitudes towards saving, risk-taking or money spending, which could help the expert in proposing improvements regarding these habits or warn of typical mistakes. This, however, also entails that experts cannot purely rely on the supposed epistemic advantages of economics contrary to other social sciences. (see Schuchardt et al. 2007, 63–65)

Concluding this chapter, testimonies from social sciences like economics, and by extension personal finance, are intimately consequential to our lives, which does not only attract virtuous attention but also frequently epistemically problematic interests. Moreover, the normativity and situatedness of social sciences render their typical testimonies more potentially ambiguous, since different perspectives cannot easily be reduced to general laws, which widens the spectrum of, *prima facie*, reasonable disagreements. Knowledge in the social sciences also involves more uncertainty due to the complex and dynamic interrelations of social phenomena. It is therefore more difficult to distinguish relevant causal factors from incidental ones, increasing the intrinsic uncertainty immanent in testimony. Economics, however, seems to have remarkably homogenous methodologies relative to other social sciences but still less compared to natural sciences. Unfortunately, it is controversial whether this is a sign of more scientific maturity. Since personal finance not only depends on economics but also psychology and sociology, any supposed advantages could not be transferred completely anyway, hence laypeople should expect a good amount of ambiguity, intrinsic uncertainty, disagreements, and detrimental interests.

In Chapter 5.1, I will discuss which consequences such a social scientific background has on the effectiveness of heuristics. For the present, I want to argue that the social scientific character makes personal finance a wicked learning environment. The distinction between kind and wicked learning environments originates from the cognitive psychologist Robin Hogarth and

describes how prevailing feedback mechanisms affect how agents can make progress in an expertise. Kind learning environments are characterised by plenty of accurate feedback, which unambiguously links to the appropriate reactions or judgments. Here, feedback is complete and straightforward, and even tacitly received information leads to appropriate inferences. Playing the violin, chess or football, for example, are expertises whose feedback allows learners to make substantive, continuous and salient progress. (see Hogarth, Lejarraga, and Soyer 2015; Watson 2020, 126–28)

Contrarily, observations in wicked learning environments do not point straightforwardly to an appropriate action or judgment. Here, feedback is often sparse or misleading. More so, mistaken beliefs can even lead to counterproductive inferences such as self-fulfilling prophecies. Watson further cites Hogarth, when he notes that while the implicit, subconscious epistemologies of humans are well adapted to kind environments, they systematically fail in wicked environments. Most expertises, unfortunately, are on the wicked side such as teaching, internal medicine or personal finance. Remembering the intrinsic uncertainty of social science, data is not easily available and rarely straightforwardly interpretable. Personal finance is by implication predisposed to render the cultivation of expertise cumbersome and limited. (see Hogarth, Lejarraga, and Soyer 2015; Watson 2020, 126–28)

4.3 Economic forecasting in personal finance

The considerations in the last Chapter addressed the social scientific background of personal finance and how this modifies the epistemic authority of putative experts. Now, I want to proceed with the special issue of (personal) finance expert practices which aim to predict the future or depend on such forecasts. More precisely, (personal) finance experts, or the specialists they, in turn, draw from, attempt to predict the future performance of financial products via analysing general market developments, potential productivity of firms or other contingencies. Such experts' effectiveness in such endeavours has been famously criticised by the studies of Philip Tetlock. His critique, however, does not only apply to financial forecasting but any long-term predictions of social scientific phenomena, such as election forecasts, for example. His extensive, long-time empirical studies showed that expert forecasters were by and large hardly able to predict events better than chance beyond roughly 1-2 years in the future. They managed to slightly outperform random samples of the population but were as good as dilettantes, who

followed world politics. Moreover, they were significantly outperformed by statistical extrapolation tools when applicable. (see Tetlock 2017, 25–66; Watson 2020, 129–32)

Besides predictions being too inaccurate, putative experts also were substantially overconfident in their judgments. The latter has been measured by comparing their proclaimed confidence in their predictive accuracy with their actual accuracy. This aspect has found support in the review of Erik Angner, who investigated Economists in particular. He found systematic overconfidence present both in lay populations and in well-regarded experts. While predictive inaccuracies can be ‘defended’ post-hoc, the same cannot be said for overconfidence. After all, if their competence allows them to predict future events, then, they also need to be able to gauge how likely it is that extraordinary circumstances meddle with the outcome, i.e. to calibrate their testimony properly. (see Angner 2006, 2–5; Tetlock 2017, 25–66; Watson 2020, 129–32)

More precisely, the more difficult the judgment tasks were, and the more confident experts were in their judgment, the more ‘overconfident’ economists had been in the end. Economic judgments, but especially forecasting, can be considered difficult due to its wicked learning environment which is empirically demonstrated by the low accuracy reached by putative experts in the field. Interestingly, empirical evidence could not find clear evidence that overconfidence increases when experts receive more relevant information or when experts have been better incentivised to testify honestly. This suggests that overconfidence in judgments in social scientific forecasting might be grounded in shortcomings of the human mind. In this sense, Angner contends: “Economists-as-experts are overconfident, I would argue, not because they are different from everyone else, but because they are just like everyone else” (2006, 7). The availability of feedback is especially bad, due to the long-time horizons until the prognosis can be compared to what actually came to be. After all, the complexity of social phenomena renders causal explanations even more difficult since effects can have multiple dynamic interrelations with other factors over this extended time. This could also explain how predictions in expertises like poker and bridge playing are substantially more precise and better calibrated. Also, meteorologists managed to create extraordinarily precise and calibrated forecasts in comparison, even if weather systems involve a good amount of intrinsic uncertainty themselves. (see Angner 2006, 4–10; Hogarth, Lejarraga, and Soyer 2015; Lane 2014; Tetlock 2017, xvii–xix; Watson 2020, 126–32)

Such environments also make it easy for putative experts to avoid realising their mistakes by giving post-hoc rationalisations for deviations from their predictions. Tetlock, for instance, mentions the habit of ‘off-on-timing’, i.e., ‘the outcome I predicted has not happened yet, but

it will!’ Another popular move is invoking ‘close-call counterfactuals’, i.e., ‘the outcome I predicted would have happened but for an incidental exogenous shock.’ The problem of forecasting social phenomena is that such excuses are not absurd. Self-immunization might, however, prevent agents from reflecting on their methods and improve upon them. Moreover, Tetlock and Angner bemoan that expert forecasts are often coated in ambiguity, for instance by omitting probability measurements, allowing ‘wiggle room’ to post-hoc frame one’s prediction as a success. (see Angner 2006, 2–5; Tetlock 2017, 121–43; Watson 2020, 126–32)

The humbling insights so far notwithstanding, Tetlock’s results have been more subtle. Contrary to his intentions, his work has been misinterpreted by parts of the political landscape as a sweeping rejection of the value of expertise. Besides forecasting not being the only cognitive expert practice, his data shows that a subgroup of the sample has been, significantly more precise and well-calibrated compared to their colleagues and compared to chance, even if not by a large degree. This superior performance did not correlate with political ideology, professional background, or reputation but with the style of reasoning. Accordingly, Tetlock distinguished two types of thinking styles, which he called hedgehogs and foxes.⁴¹ A typical hedgehog is a big-picture theorist who is well-versed in one theoretical framework, which she uses to interpret data and confidently predict future developments. Foxes instead are more open-minded in comparing multiple different methods and theories and more humbly anticipate alternative possibilities. It has been the latter whose complex methodologies significantly outperformed the confident and simplistic hedgehogs both in accuracy and calibration. Hedgehogs were too intellectually invested in their theories to recognize their limits and alternatives. They were also more prone to engage in post-hoc rationalisations.⁴² (see Tetlock 2017, xvii–xix, 67–120; Watson 2020, 129–32)

41 Tetlock adopted the metaphor of the fox, who knows many little things and the hedgehog, who knows one big thing from Isaiah Berlin’s “The Hedgehog and the Fox” (see 1953). According to Fuller and Tetlock himself, this metaphor was first popularised by Erasmus of Rotterdam but could possibly be traced back to the Ancient Greek poet Archilochos. The long history notwithstanding, Tetlock uses this metaphor freely to distinguish thinking styles in putative forecasting experts. (see Fuller 2018, 251; Tetlock 2017, xx)

42 Following up on these empirical findings, Tetlock and his colleagues developed the ‘Good Judgement Project’ where they tried to better delineate the ‘foxy’ way of thinking that enables the most accurate predictions by organising forecasting tournaments. In his book *Superforecasting* with Dan Gardner, they present their first lessons on how to improve forecasting. For instance, they were able to train teams of forecasters such that they followed norms of open-minded exchange, and outperformed people working alone. Remembering the belief-independence thesis, this empirical insight suggests that agents can cooperate productively without being too epistemically dependent on each other. Moreover, Tetlock shows that in aggregation crowds with multiple different knowledge backgrounds perform better than like-minded individuals. The latter gives empirical support to robustness intuition, at least in an aggregational sense. Nevertheless, predicting the future in complex contexts remains an uncertain task. (see A. I. Goldman 2018, 7; see Tetlock 2017, xxvi–xxvii; Tetlock and Gardner 2015)

The foxes' self-critical, point-counterpoint style of thinking prevented them from building up the sorts of excessive enthusiasm for their predictions that hedgehogs, especially well-informed ones, displayed for theirs. Foxes were more sensitive to how contradictory forces can yield stable equilibria and, as a result, "overpredicted" fewer departures, good or bad, from the status quo. (Tetlock 2017, 21)

In this chapter, I discussed expert forecasting as part of a (personal) finance expertise. Unfortunately, it seems that economic forecasting is too imprecise to give recipients an epistemic advantage. Returning to my distinction between basic and forecasting expertise in personal finance, I argue that the former is subject to low to moderate intrinsic uncertainty due to its social scientific background but otherwise moderate scope. Forecasting practices on the other hand are subject to very high intrinsic uncertainty, rendering their epistemic benefit meagre. In Chapter 5.2, I will address which consequences this inability has on heuristics to assess personal-finance experts.

5. Applying the general debate to the domain of personal finance

5.1 Assessing incentive structures in personal finance

In Chapter 3.1, I discussed the general expert identification debate regarding biases and detrimental incentives, i.e. its honesty dimension, and complemented it with a strategic, game-theoretic perspective. By analysing cheap talk models, I showed with formal support that experts whose interests are misaligned with laypeople can abuse the existing information asymmetry to give misleadingly precise advice. *Prima facie*, laypeople should anticipate the bias and only believe in the more general thrust of such testimony proportional to the size of the bias, unless interests are diverging so far that the testimony becomes fully uninformative. Even though many scholars of the philosophical debate were aware that biased expert testimony should be avoided, they, except for Almassi and Duijf, remained rather vague about which reactions are warranted for what magnitude of interest misalignment. I concurred with Duijf and the cheap talk literature that moderately biased testimonies can still be informative. Instead of missing relevant information, laypeople can justifiably believe in a more ambiguous version of the testimony. (see Almassi 2017; Crawford and Sobel 1982; Duijf 2021; A. I. Goldman 2001, 104–5; Grundmann 2024, 10–12)

To follow up on this more nuanced investigation, I introduced Almassi's conceptual analysis of detrimentally incentivised testimonies to analyse what conditions constitute them.

Accordingly, the further misaligned the preferences between the parties are, the more the propositions allow potential ambiguity, and the more practical opportunity experts have for acting on their interests, the stronger biased laypeople should expect testimonies to be. Understanding these conditions in a specific context helps laypeople detect detrimental incentives and gauge how far they could avoid imprecise testimony. This is of high heuristic value since detrimental incentives are social facts, accessible to laypeople and thereby one of the few possibilities to make direct heuristic judgments about expertise and rationally ground trust for other heuristics. The question remains, however, how this abstract debate on incentive structures fares in the context of personal-finance expertise. (see Almassi 2017)

5.1.1 The ambiguity and uncertainty of personal-finance advice

I begin by revisiting the potential ambiguity criterion. The more propositions allow for a larger range of possible answers the more theoretical opportunities experts have for giving a misleadingly precise answer, which ought to be rationally discounted into a more ambiguous message by laypeople who (imperfectly) anticipate the bias. I will argue that basic personal-finance expertise as a social science features a heightened level of potential ambiguity in its propositions by drawing from insights from Chapter 4.2. (see Almassi 2017; Grable and Chatterjee 2022b)

Firstly, as I have argued, social sciences do not just study impersonal causal mechanisms but also aim to understand the intentions and reasons of agents. Their research objects are social phenomena that feature a normative element since scientists need to distinguish between properly rational and mistaken actions. This situates explanations of social interactions into multiple perspectives which can be investigated by scientists, introducing a subjective-perspectival element. Consequently, multiple, *prima facie*, reasonable explanations for the same problem can exist, which simultaneously increases the potential ambiguity of answers. For example, both mental overload or carelessness could explain the hesitancy of an agent to meticulously review her debts, which again could be interpreted either as rational energy-saving or myopic discounting of future well-being. Hence, these normative and perspectival elements allow testimony to become more ambiguous. (see Almassi 2017; Grable and Chatterjee 2022b; Hausman 2021, chap. 2.2-2.3; Reiss 2014, 139–40)

Moreover, social phenomena are embedded in a very complex interplay of (potentially) dynamic causal interactions. Analysing the interactions of many agents involves various

possible causalities, which warrants decisions regarding which factors explain a phenomenon. Especially, the possibly dynamic interactions of factors and their potentially uncertain frequencies are messy and difficult to assess. This causal complexity results in propositions, which can have a wide range of answers adding a further reason why typical testimonies in personal finance allow for a heightened amount of ambiguity. (see Almassi 2017; Boumans and Martini 2014, 1–4; Hausman 2021, chap. 2.3; Lane 2014)

Turning to the practice of economic forecasting in personal finance, testimony here is often issued intentionally vague by putative experts. Tetlock and Angner both bemoan that forecaster often omit to (precisely) qualify time, intensity, probability and confidence in their predictions. This allows them to feign precision while their testimony is often very ambiguous. I will again address the tactical reasons and consequences of this in Chapter 5.1.3. Making predictions about the developments of social phenomena years into the future that were already complex in the presence increases the intrinsic uncertainty immensely. Given the inaccuracies Tetlock and Angner discovered, any less ambiguous testimonies would be unrealistic anyway at the current state of the art. Compensating for this intrinsic uncertainty, questions are cast with high potential ambiguity. (see Angner 2006; Lane 2014; Tetlock 2017)

5.1.2 Common incentives in personal finance

In Chapter 4.2, I argued that personal finance, like other social sciences, involves perspectival and normative elements due to the social phenomena they study and the human interaction they want to explain. Though such methods might be necessary to fully understand human behaviour and draw advice from it, this subjectivity might also increase the influence of value judgments on descriptive propositions. Following Reiss, since some epistemic norms are necessary for doing research, it might not be easy to uncontroversially determine which values are important, harmless or instead biasing. Without intending to digress into the philosophical debate on values in science, I want to suggest that personal finance, due to its difficult discriminability regarding value involvement, might be more prone to biases. (see Hausman 2021, chap. 2.2-2.3; Reiss and Sprenger 2020, chaps 2–3; Reiss 2014, 139-141, 149-151)

In Chapter 4, I further contended that personal-finance issues could be of more intimate concern to laypeople. Contrary to many political issues, here, a person's decision affects her well-being in a more immediate, visible fashion which might increase her motivation to address these issues. Following Hausman, I further argued that economics, and by implication personal

finance, is not just intimately connected with the interests of the laypeople but also to those of the putative experts. Often testimonies in these disciplines can provoke policies that are lucrative to some parties independent of their general truth. For example, a proposition for tax law might transparently address a goal like ‘generating the national budget most efficiently without burdening citizens too much’. Different ways to implement taxation, deriving from different reasonable theoretical backgrounds might impact various social group’s interests differently and thus the choice of taxation can easily be affected by the idiosyncratic interests of the expert’s social group. The same is true for personal finance, as the decisions of clients can easily benefit the expert independent of the former’s best interest. Experts can, for example, urge laypeople into buying financial products they receive a commission from, into spending more money than they need or into pushing a stock from which they benefit. Accordingly, this suggests that we should expect biased interests more frequently in the context of personal finance. (see Hausman 2021, chap. 2.2-2.3; Reiss and Sprenger 2020, chaps 2–3; Reiss 2014, 139-141, 149-151)

To what degree these two arguments drive the occurrence of biases can, nevertheless, be questioned. Interestingly, Miller divided the disciplines into social or special sciences including biology on the one side and physics and mathematics on the other side, where the latter is not even an empirical discipline. Indeed, climate sciences, biology or medicine all belong to the natural science but still feature frequently in the expert-identification debate as examples where detrimental interests impede the trustworthiness of testimony. The Wakefield MMR vaccine case, the denial of anthropogenic climate change and the harmfulness of tobacco smoke, here, are all prominently discussed examples which are based on natural sciences. Accordingly, the occurrence of detrimental interests might be more driven by the relevance and consequences for individuals and the public discourse rather than the influence of values in social sciences. Since personal finance interacts with interest in both ways, laypeople should be vigilant and expect a heightened occurrence of biases anyway. (see Anderson 2011, 149–53; Collins and Weinel 2011, 404–11; Evans 2014, 241–45; A. I. Goldman 2021, 92–94; Miller 2013, 1313–14; Oreskes 2012)

Despite recent examples, not all detrimental incentives in personal finance are solely financial. De Bruin, for example, discusses epistemic injustice in finance and cites empirical literature on how female borrowers have to pay higher interest rates on their loans with less favourable conditions, are more frequently rejected loans entirely or face less favourable treatment when defaulting on one. If such treatment is caused by unjustified prejudices or

systematic issues, then these are examples of biases not caused by monetary interests but (sexist) moral beliefs. While biases regarding social identities might be more effectively combatted by collective or political action, this would go beyond the individual scope of my thesis. (see Almassi 2017; de Bruin 2021)

I will continue by discussing a standard example of how detrimental incentives can affect personal-finance advice, which has been presented by Inderst and Ottaviani. They analysed the practice of commission-taking in the context of experts selling recommendations for suitable financial products to laypeople. For example, investment funds, stocks, currencies or insurance. In this context, this means an adviser receives remuneration from a (financial) product provider for inducing the customer into buying from the provider's assortment. Since the provider, of course, profits from any product they sell they have a reason to incentivise advisers to compel their clients into buying their product. The adviser, on the other hand, can receive additional rewards from recommending the provider's products through this commission compared to alternatives. This gives the adviser a reason to recommend an option independent of the best interest of her client. This dependent expert advice could take the form of a bank clerk deciding on recommending 'in-house' products, which are provided by an investment firm owned or affiliated with the bank. Nevertheless, there are many varieties of how the basic rationale just described presents itself in practice. Notably, Goldman mentioned this issue as an example of the effects of detrimental incentives on trustworthiness in his 2021 paper: (see Inderst and Ottaviani 2012)

An important additional factor to consider, of course, is the self-interest an expert might have in recommending one way of addressing a problem rather than another, for instance by getting 'kickbacks' from choosing one investment outfit rather than another. (A. I. Goldman 2021, 91 emphases in original)

Note that the expert advice in this context does not need to substantially depend on economic forecasting, since Inderst and Ottaviani focus on basic personal-finance expertise as characterised in Chapter 4.1. Accordingly, their contribution is mostly independent of the problems of forecasting. Moreover, their analysis focuses on a central aspect of personal-finance expertise but does not cover the entire range of testimony the expertise can offer. (see Grable and Chatterjee 2022b; Inderst and Ottaviani 2012; Schuchardt et al. 2007)

I will shortly outline their economic model and discuss the incentive structure they describe. While their paper focuses on policy proposals, I will investigate consequences for individual heuristics. To analyse the effects of payment modalities, they use a cheap talk model as I have

discussed in Chapters 3.2 and 3.3. However, they extend it by an additional step prior to the expert advising her client, where the expert is contracting with a monopolist financial product provider, i.e. an investment firm. The layperson has the choice between buying an advanced product for a price (management fee) issued by the provider and the outside option, i.e. a competitive standard product or nothing at all. While she does not know which option is more suitable for her, the authors assume that the outside option is more *likely* suitable. (see Inderst and Ottaviani 2012, 394–401)

Note, that the equilibrium in this model, though following similar rationales, does not fully resemble that of the original cheap talk models after Crawford & Sobel, since, firstly, it is extended by an additional contracting level and, secondly, the client has two discrete options which following the adviser's testimony. After contracting an absolute and a product-conditional compensation with the provider, which the client does not know, the expert learns about the option's suitability to the client and engages in cheap talk with the latter. Contrary to the provider, the adviser also partially cares about the satisfaction of the client, which could be understood as the ability of the latter to somewhat punish the expert after realising low utility. Moreover, the adviser can demand a flat advising fee from the layperson. In the course of their analysis, Inderst and Ottaviani distinguish two cases: One with wary clients who perfectly anticipate the consequences of commissions and one with naïve clients who do not do so, whereas the provider and advisor know what kind of client they face. (see Crawford and Sobel 1982; Inderst and Ottaviani 2012, 394–401)

In the case of wary clients, the product provider can anticipate that the adviser will set the advising fee to the highest acceptable limit and can in turn set the absolute compensation as high as the indifference level of the adviser, which is the negative of the entire advising fee.⁴³ Thereby the monopolist provider can reap the surplus of both the client and the adviser and set the price to 0 to maximize joint profits. In this equilibrium, the product provider reaps the profit, the adviser has no reason to act on anything but the suitability to the client and the latter pays an advising fee by its market price. Since the provider could anticipate that the wary customer would anticipate that a pricing shift from the advertising fee to the advanced product price would not amount to less costs but make it beneficial for the provider to commission the adviser,

43 The possibility of setting this value negative, allowing the provider to extract rather than pay compensation is a conventional modelling idealisation. This resembles contractual flexibility where the provider and adviser are closely connected. In a later chapter, the authors analyse the situation where this absolute compensation cannot be negative. This has, however, hardly any consequences for the main insights of the paper besides that when contracts are made at arm's length, then testimonies for wary clients are biased a bit. (see Inderst and Ottaviani 2012, 403)

neither agent has a reason to deviate from this setting. (see Inderst and Ottaviani 2012, 396–401)

Naïve clients, however, would not understand that a reduction in advising fee would, in equilibrium, be optimally added to the price of the advanced product and induce commission-taking of the adviser. Anticipating this, the adviser and, in turn, the provider can profit by reducing the up-front advising fee to induce more trade, meanwhile increasing the product price. Moreover, the provider can increase her surplus by incentivising the adviser to recommend advanced products more often than it would be suitable for the layperson. The naïve client underestimates the likelihood that a product is being recommended by the adviser when it is not suitable. Hence, the naïve customer would, in equilibrium, pay no up-front fee but receive biased advice and a generally lower expected utility. (see Inderst and Ottaviani 2012, 396–401)

So far, Inderst and Ottaviani have formally shown that when laypeople do not adequately anticipate the effects of commission-taking on financial products, they pay less advising fee up-front but instead pay relatively more for higher product prices due to less suitable advice. The better laypeople, however, understand this mechanism the more they need to pay an up-front advising fee yet receive recommendations for cheaper and more suitable products which makes them comparatively better-off. As Inderst and Ottaviani say, this might explain why dependent expert advice, for example at one's house bank, is commonly 'free'. (see Inderst and Ottaviani 2012, 396–401)

Inverting the question, I want to now discuss whether laypeople could ever be epistemically justified in preferring advice from dependent experts. Having Inderst and Ottaviani's analysis in mind, the answer seems to be *no*, but they also discuss a potential exception. When the expert's effort of information collection is included in the model, commissions could properly incentivise additional energy investments into finding highly specialised products for wary clients. By implication, a ban on dependent advice could eliminate this opportunity. Without going into detail, I want to argue that this exception is hardly relevant since its premises are unlikely. (see Inderst and Ottaviani 2012, 403–5)

Firstly, hardly any layperson, as I have described in Chapter 2.1, could be 'wary' in Inderst and Ottaviani's sense. Empirical studies and surveys in the past have shown that most clients were not fully aware when personal-finance experts took commissions and neither understood how this biased their testimony. Inderst and Ottaviani show this by citing EU surveys, reports

and studies showing that laypeople systematically receive unsuitable advice from which their counterparts profited. Of course, this does not mean that laypeople cannot understand incentives. On the contrary, assessing the incentive structure of an expertise is an important heuristic for laypeople, which I am emphasising in my thesis. Also, laws could demand disclosure clauses in contracts which elucidate the involved relationships. (see Inderst and Ottaviani 2012, 394–96, 401–5, 407–8)

Nevertheless, I want to argue that while knowledge about the size of the bias alone is not enough to disincentivise it (one needs to be able to ‘punish’ as well), neither is it likely that laypeople are able to perfectly assess the size of the bias. Behavioural research by George Loewenstein, Cass R. Sunstein, and Russel Golman shows that disclosure policies, *simpliciter*, are not effective in inducing salient understanding in recipients. Also, Inderst and Ottaviani cite evidence in experimental cheap talk studies that laypeople trust advice more than they should. Unfortunately, financial product providers and experts do not have an interest in improving the quality of their disclosure more than they are legally required to do, since they always profit more from naïve than wary customers. Also, Angner referred to empirical research that laypeople are not able to adequately discount testimony they expect to be overconfident or biased since people are also influenced by information they do not want to believe and neither feedback nor repetition helps in this respect. Furthermore, as analysed in cheap talk models, even if the bias size were known, testimony could be ambiguous which disables said discounting strategies. In conclusion, I argue that laypeople with constraints in money, time, and education are hardly able to perfectly anticipate biases. This shows that laypeople using heuristics to detect detrimental incentive structures might only be able to approximate the true size. This impedes their ability to adequately assess the honesty and thus trustworthiness of putative experts in this respect. While this demonstrates how the imperfections in assessing biases are likely to persist in the context of personal-finance advice, it does not show that this discipline is more or less afflicted by this problem than other expertises. (see Angner 2006, 15–17; Inderst and Ottaviani 2012, 394–96, 401–5, 407–8; Loewenstein, Sunstein, and Golman 2014)

Secondly, I find it puzzling to imagine which laypeople could profit from highly specialised financial products. I think it sensible to expect most laypeople to have only moderate financial means which reduces their scope. Moreover, long-time forecasting is hardly possible as Tetlock and Angner have shown in Chapter 4.3. Therefore, I expect the margin of laypeople who could profit from increased efforts in information collection who are not already professional

customers to be very small if not non-existent. Thirdly, the product provider also needs to supply a wide range of products for commission-taking experts to have something to search for, yet especially dependent advisers are incentivised to a constrained assortment of products the provider wants to sell. (see Angner 2006; see Inderst and Ottaviani 2012, 403–5; Tetlock 2017)

Consequently, it is virtually never sensible for laypeople to prefer dependent advice when independent testimony is available, since at the likely best (them being able to perfectly anticipate and reduce the effects of commissions) it is equivalent to independent advice. More likely, however, they will be worse off as a result. Since the ability to assess incentives and combat them is limited, the best option seems to rather avoid them. Moreover, if advice from an expert is free then a layperson should be vigilant and question how else the expert is financing herself and how this affects the advice. (see Inderst and Ottaviani 2012)

5.1.3 Opportunities to bias personal-finance advice

In Chapter 3.1, I discussed Almassi's opportunity criterion, which addresses how practical opportunities for experts to bias their testimonies affect their trustworthiness. In this vein, I discussed legal constraints and institutional control of expert testimonies which are designed to disincentivise detrimental interests. Moreover, I discussed reputation mechanisms and individual bargaining power, i.e. possibilities to punish or avoid badly incentivised experts so that biased experts would regret it. Now I am turning to the epistemic environment of personal finance to discuss the extent of their practical opportunities to indulge their detrimental interests. (see Almassi 2017)

In Chapter 5.1.1, I argued that typical testimonies in personal finance feature a higher potential ambiguity and higher intrinsic uncertainty. The practical consequence of this issue is that Goodwin's rationale of an 'appeal to expert authority' is significantly weakened. His reputation mechanism depends on the idea that experts openly position themselves and consequently can be trusted on these testimonies since they would be punished if laypeople realised their bias over time. The point of concern is whether laypeople can reliably enough recognize biases post-hoc and punish accordingly. Remembering Chapter 2.2, I already argued that laypeople would not reliably be able to distinguish detrimental interests from intrinsic uncertainty as the source of expert disagreement. For the same reason, laypeople cannot discriminate post-hoc between bad performance due to biases or uncertainty intrinsic to the testimony. The higher the intrinsic uncertainty the less accurate are testimonies which renders

discrimination increasingly difficult, yet blanket punishments would also affect humble but honest advice. Testimony with high potential ambiguity on the other hand would allow experts to avoid detection of biases or bad performance since it allows them to post-hoc reinterpret their testimony within the wider range of possible states of their testimony. I argued in Chapters 4.2 and 5.1, that personal finance is prone to heightened potential ambiguity and intrinsic uncertainty, which impedes Goodwin's mechanism of reputation and thus increases the putative expert's opportunity to bias. (see Almassi 2017; Goodwin 2011, 288–94; Lane 2014, 105–11)

This issue becomes more pronounced when we turn to the practice of (personal) finance forecasts. Here, I argued in Chapter 4.3, that intrinsic uncertainty is high enough to render any testimonial knowledge transmission questionable. Accordingly, Tetlock and Angner criticise experts' conduct of post-hoc reinterpreting their strategically ambiguous testimonies when they fail to be accurate. Since such post-hoc rationalisations are not always absurd, as I have described in Chapter 4.3, detecting and punishing either bad performance or biases is difficult. Tetlock and Angner, moreover, suggest that being overconfident regarding intrinsic uncertainty in fields like personal finance could be strategically advantageous for putative experts. In this vein, Tetlock argued that “the reputational benefits of standing by outlier predictions that eventually prove right far exceed the penalties that attach to being defiantly wrong in relative obscurity.” (2017, xxxv) (see Almassi 2017; Angner 2006, 10–14; Goodwin 2011, 288–94; Tetlock 2017, xxxv–xxxvi)

Respectively, Tetlock observes that hedgehog-type experts receive more attention in the media even though they are less accurate. Steve Fuller adds the interpretation that such hedgehogs should rather be called lions, inspired by Machiavelli, as they aim to prescribe the world as it *should* come to be instead of predicting what is *likely*. The ability to speak about future events in memorable and captivating narratives might additionally explain their comparative popularity in the media. Specifically concerning their overconfidence, Angner adds that a selection bias might further help in understanding this phenomenon since very confident agents might also be more likely to offer their service and push into the public. Moreover, people seem to have psychological tendencies to interpret confidence as an indicator of competence more than it is prudent. Angner concludes that economists do not face effective social sanctions to reduce ambiguity or overconfidence in their testimonies. (see Almassi 2017; Angner 2006, 10–14; Fuller 2018, 250–52; Goodwin 2011, 288–94; Tetlock 2017, xxxv–xxxvi)

Regarding the problems laypeople have in discriminating bad and biased expert performance, seeing through post-hoc rationalisations, or understanding intrinsic uncertainty in

personal finance, they could profit from listening to meta-expert testimony. As I have argued in Chapter 3.1, if laypeople can justifiably trust these, they can use their expert testimony to gain a more nuanced understanding of the trustworthiness of putative experts. Since the practical personal-finance advice I am analysing in this paper is not scientific itself (though it is informed by several scientific disciplines), their specific testimony is not checked by a scientific institutional framework. I am emphasising this point because the general debate on the expert-identification primarily discussed testimonies from the (natural) sciences and, as I have suggested, implicitly adapted their debate to this context. Especially scholars like Anderson and Grundmann have emphasised the importance of scientific institutions in justifying the credibility of (scientific) expert testimony. Accordingly, personal finance testimony misses the meta-expert assessment involved in the degrees, grants and prizes issued by universities, peer review mechanisms or the objectifying practices in scientific methods like norms of representative data collection or ‘blinding’ in RCTs. The question of interest regarding personal finance is, then, whether experts here have fewer institutional constraints to bias testimony or whether other frameworks can compensate for this. (see Anderson 2011; Grundmann 2024, 15–17)

Alternatively, the discipline of personal finance knows private, non-profit standard boards whose mission is to define good and integer practice. In addition, they examine and certify putative experts, which renders them a prime example for a meta-expert institution. If they can rationally instil trust in laypeople due to, for instance, their good reputation, the latter can justifiably increase trust in the bearers of these certifications. The most popular example is the Financial Planning Standards Board (FPSB), which is an umbrella organisation encompassing national boards in many nation-states. Associated organisations like the ‘Österreichischer Verband Financial Planners’ issue certificates such as the CFP (Certified Financial Planner) from the FPSB or the EFA (European Financial Advisor). (see ‘FPSB’, n.d.; ‘About FPSB | FPSB’, n.d.; ‘Verband Financial Planners’, n.d.; ‘Verband’, n.d.; ‘Karriere’, n.d.)

Finance-related degrees of higher education could be an indicator of competence as well, as universities and colleges function as meta-experts, albeit less than in the context of scientific expertise. While such degrees indicate knowledge of financial mechanisms, they do not indicate advisory skills. Previous (work) experience in a related field has also occasionally been proposed as an indicator of competence in the expert-identification debate but I deem it difficult for laypeople to understand which specific career backgrounds are reputable and relevant within

the community. Also, experience in a CV does not mean the putative expert has done a good or integer job. (see A. I. Goldman 2001, 97; Goodwin 2011, 288–90)

Other control institutions relevant to personal-finance issues are public banking and finance regulation authorities like the Financial Market Authority (FMA) of Austria. The FMA is an Austrian state bureau tasked with supervising the national financial market and thereby enjoys extended autonomies. It supervises credit institutions, investment firms, and their product information documents, issues punishments for misconduct and warns laypeople of malicious actors. What norms their supervised institutions need to follow and what constitutes misconduct is defined by the legal frame of the jurisdiction, which in this case is an Austrian or EU context. (see ‘FMA Austria’, n.d.; ‘Financial Market Supervision in Austria’, n.d.; ‘Tasks of the FMA’, n.d.; ‘European Supervisory Architecture’, n.d.)

Relatedly, the legal framework changes the incentive structure and thus the trustworthiness of putative personal-finance experts in a jurisdiction. In the Austrian example, the MiFID-2 directive of the EU should be mentioned. This directive has been transposed into Austrian law via the Securities Supervision Act 2018 (Wertpapieraufsichtsgesetz, WAG 2018). Financial product providers now are obliged to define a target audience and transparently explain how the risk/performance profile, the recommended holding period and other criteria are tailored accordingly. Moreover, credit institutions, investment firms and advising agencies need to protocol personal meetings with their clients, record telephonic and electronic communication and store these records for at least 5 years. The intention behind the latter is to render false advice and fraud verifiable.⁴⁴ (see Beek 2022; Inderst and Ottaviani 2012) (see Directive 2014/65/EU, WAG 2018)

When taking personal-finance advice, the client must receive a suitability statement containing what recommendations were given and how they are supposed to fit the client’s needs and preferences before any order is executed. The client also needs to receive a report about the fees of the product and the advising service, which allows them to transparently compare the costs and returns of the product. This legal framework also defines independent and non-independent investment advice. Independent personal-finance experts need to evaluate a broad range of potentially suitable products from different providers and are not allowed to take commissions or have any other economic associations with such firms. Finally, this

⁴⁴ Speculatively, this could be the reason why nowadays most consultations are happening locally on an eye-to-eye basis, since then at least actual wordings cannot be held against the (putative) experts. Alternatively, such a mode might just be more effective for invoking customer loyalty.

framework also sets requirements for the quality and training of personal-finance experts, which ought to be incorporated into certified education programs and are controlled by, amongst others, the Austrian Financial Market Authority (FMA). (see ‘FMA Austria’, n.d.; ‘Financial Market Supervision in Austria’, n.d.; ‘Tasks of the FMA’, n.d.) (see Directive 2014/65/EU, WAG 2018)

Reviewing these meta-expert institutions and laws, it becomes clear that personal-finance advice is not happening in an institutional vacuum. While I suggest that the institutionalisation of scientific knowledge generation with its pervasive expert peer control, transparency measures and methods whose main goal is to ensure safe knowledge might be, *ceteris paribus*, more powerful in licensing trust, I think personal finance is also substantially regulated. Indeed, given the high ambiguity and intrinsic uncertainty of personal finance along with the pervasiveness of corrupting incentives, the legal consumer protection pulls a lot of weight in incentivising trustworthiness. Whereas the certificates of trusted standard boards can indicate competence. The knowledge about legal restraints and threats of punishment by misconduct justifiably increases the general trustworthiness of personal-finance experts. Accordingly, laypeople benefit from cursorily checking which kinds of laws are present and how they are being executed to calibrate their trust in personal-finance experts within their jurisdiction. (see Directive 2014/65/EU, WAG 2018)

The fact that commission-taking while advising laypeople, however, is not outlawed in most of Continental Europe is a noticeable disadvantage considering my previous discussions. Also, disclosure policies might not raise sufficient awareness in laypeople, showing the limits of such transparency policies and controls to disincentivise biases. In practice, of course, the threat of punishment also depends on the ability of these control institutions to meaningfully execute these punishments. As argued in Chapter 3.1, the FMA might be able to enforce transparency and punish fraud but is not able to improve the honesty of advice beyond a threshold of plausible deniability and suitable *enough* advice. After all, assessing each or even just a sample of advice on their best possible quality is almost as much work as giving the advice oneself and thereby unfeasible for control institutions. Unfortunately, considering the high level of intrinsic uncertainty and heightened potential ambiguity of the controlled testimony, the threshold of plausible deniability is inconveniently low. (see Beek 2022; Inderst and Ottaviani 2012, 401–2, 407–8; Loewenstein, Sunstein, and Golman 2014) (‘FMA Austria’, n.d.; ‘Financial Market Supervision in Austria’, n.d.; ‘Tasks of the FMA’, n.d.)

5.1.4 A further heuristic utilizing strategic thinking

Apropos strategic thinking in heuristics, I want to propose a further rationale that utilizes this dynamic interaction of intentions between agents in a similar vein to Goldberg's argument described at the end of Chapter 2.2. Laypeople confronted with testimony of the kind 'follow my strategy and become rich' or 'invest in these products and receive high returns with low risks' could question the consequences this testimony had were it true. If the testifier knew a strategy to become rich, this knowledge entails in the context of personal finance that she must have observed some agents who used this strategy and became rich because of it. If this is the case and if she testifies this publicly, there need to be some people who know about this. If some people know about this and are open to sharing these insights with strangers or a wider audience, why is this news not travelling faster, given that such information is extremely demanded and newsworthy? Why then are not many more people already rich? (see de Bruin et al. 2023, chap. 4.2.2; Goldberg 2011)

Alternatively, if these are secret insights shared with you how could she profit by this over just using it herself? This leads to the second aspect specifically relating to financial products like stocks, whose value depends on the subjective evaluations of many market participants. *Ceteris paribus* and simplified, the price of, for instance, a share of a company depends on its current and expected future (trading) value, which is the result of subjective evaluations of market observers. If someone, for instance, knows more about a higher future value of a firm than other market participants, this person can use her information advantage to buy cheaper shares of the currently undervalued firm and reap future rewards. Only such information advantages over other participants can (besides the actual ability to buy products) enable profits without risk. Holding periods aside, the higher return of a product is otherwise being compensated by its heightened volatility, while the expected value theoretically stays the same. Having this in mind, it is questionable why a putative personal-finance expert is sharing such an information advantage, since if too many know about it, it ceases to be a comparative advantage, which defeats the purpose. Moreover, in Chapter 4.3 I have shown how difficult it is to predict the future of social phenomena like the stock market. This begs the question of how this individual could have attained such esoteric knowledge. (see Angner 2006; de Bruin et al. 2023, chap. 4.2.2; Goldberg 2011; Tetlock 2017)

5.2 Assessing the track record of putative personal-finance experts

In this chapter, I want to discuss whether or how Goldman's 'salient-success' heuristic could help laypeople in making heuristic judgments about putative personal-finance experts' competence. Goldman deems it as the likely most helpful heuristic. Indeed, next to the assessment of incentive structures, I think this is a rare opportunity for direct heuristic judgments which is essential in rationally grounding laypeople's trust. Consequently, assessing the possibility of adaption to personal finance is important. In this vein, this section draws on my previous argumentation in Chapters 3.2 and 4. (see A. I. Goldman 2021, 89–91; 2001, 106–8)

Goldman concedes that this heuristic is not always applicable since not all expertises predict phenomena whose success condition is salient to laypeople. More so, not all kinds of expertises involve predictions but Goldman's rationale also works with other kinds of salient successful outcomes which let laypeople infer expertise as their likely cause. For example, a mechanic might more likely be an expert if she seemingly returned function to a broken car. Still, Goldman further concedes that there are several contexts and disciplines where this approach is not working. Interestingly, the following quote from his 2021 paper seems to suggest that he deems investment advice as a context where the 'salient-success' heuristic is working rather well: (see A. I. Goldman 2021; 2001; Grundmann 2024)

To take simpler examples, laypersons can sometimes successfully identify cases in which their doctor, their mechanic, their investment adviser, etc. gives them just the sort of treatment, advice, etc. that palpably succeeds in generating sought-after results. (A. I. Goldman 2021, 90–91)

Locating 'investment advice' in my conceptualisation of personal-finance expertise, I think this notion is closer to the forecasting aspect than basic personal-finance expertise since the function of good investing is to make the right choices for future benefits. This mostly involves choosing the right financial products in the hope they will perform well in the future to increase one's capital or avoid loss due to inflation. *Prima facie*, one might intuitively think about investment advice as having salient success conditions. After all, we might know of or can imagine success stories of people who seemingly made a fortune in investing in the right stocks. It is a phenomenon of popular media to ask 'successful' stock traders for advice, implying that such people have special expert insights into financial markets. Also, the selling point of 'actively managed' investment funds is that experts analyse the market for you to select financial products that perform better than, for example, a fund which simply buys according to a fixed market index. Customers pay these experts a management fee to profit from their putative information advantage. This is a very common example, where forecasting expertise is being

assumed in the real world. (see Angner 2006, 7–10; de Bruin et al. 2023, chap. 4.2.2; A. I. Goldman 2021, 89–91; 2001, 106–8; Tetlock 2017, xxxv–xxxvi)

Explicating this intuition, laypeople might assume that the ‘salient-success’ condition is fulfilled when one multiplies her available funds by investing. The rationale is that this success is likely to be the result of a special understanding of the market and how products will perform in the future. Even though this seems to be a prime example of Goldman’s ‘salient-success’ heuristic, I want to argue that it is only apparently so. Firstly, consider that not everyone who engaged with financial markets and became rich did so because of one’s investments. Hasty assessments and sloppy journalism might overlook the moderate contribution one’s investments had on one’s finances. After all, the person in question might have started with an already remarkable capital or gained money by other means. Moreover, the observation of successful investors can suffer from survivorship biases. Statistically, it is likely that few market participants out of many coincidentally choose the right products over a semi-long period just as it is likely that a few will incidentally perform below average. If we only focus on the success stories, we will have a skewed perspective of reality and mistake luck for expertise. (see A. I. Goldman 2021; 2001; Grundmann 2024)

While I just presented an alternative explanation for apparent success, the empirical investigation of Tetlock and the analysis of Angner, raise doubts that financial forecasting expertise as advertised exists at all. As I already discussed in Chapter 4.3, financial predictions next to other political forecasts quickly decline in accuracy to mere chance as the focus becomes too narrow or the time horizon extends over more than one or two years. Such forecasting, however, is too myopic to assess whether a financial product will be profitable for a layperson. This might be a consequence of the complex and dynamic interrelations that affect markets and the high intrinsic uncertainty experts therefore face concerning longer periods into the future. Also, to profit from forecasting in the stock market, agents additionally need an informational advantage compared to others or the prices will already reflect the expected future value of a product. Hence, Tetlock’s thorough longtime empirical studies show that by implication salient successes in personal-finance forecasting are illusory. (see Angner 2006; de Bruin et al. 2023, chap. 4.2.2; Evans 2014; Tetlock 2017)

Returning to the example of actively managed investment funds, there is no evidence that such products perform significantly better over a long period than simple Exchange Traded Funds (ETFs) that automatically replicate stock market indices like MSCI World. More commonly they underperform such benchmarks and might be, in combination with the

management fee they require, a suboptimal choice for laypeople. These findings notwithstanding, actively managed investment funds are very common. This circumstance moved some philosophers to argue that such arrangements should be considered fraudulent. (see Angner 2006, 7–10; de Bruin et al. 2023, chap. 4.2.2; Tetlock 2017, xxxv–xxxvi)

What does this now entail for the basic part of personal-finance expertise? Such practices barely depend on specific forecasts but rather on time unspecific causal relations. Tasks like structuring one's debt, developing a simple saving plan, finding suitable insurance, or optimizing one's taxes are therefore not subject to Tetlock's critique. This aspect, unfortunately, neither has any salient markers of success accessible to laypeople. A personal finance intervention that results in well-managed finances is not obvious enough to laypeople. Differences might be observed within a precise professional analysis, but laypeople would not have the time or education for such endeavours. Moreover, being financially secure can have many reasons and vice versa events can happen that disturb even well-managed finances. After all, such social (scientific) phenomena involve heightened intrinsic uncertainty as I discussed in Chapters 4.2 and 4.3. Hence, Goldman's heuristic cannot be applied to personal-finance expertise. (see A. I. Goldman 2021; 2001; Grable and Chatterjee 2022b; Tetlock 2017)

I think my argumentation in this chapter is a good example of how a more detailed analysis can benefit the expert-identification question. Expertises differ in characteristics that are consequential for laypeople. Conversely, a more particular analysis of specific expertises can inform the general debate and better outline the conditions for and scope of heuristics. Not just personal-finance expertise but no discipline that depends on longtime social scientific forecasting has salient success conditions. (see Angner 2006; Evans 2014; A. I. Goldman 2021, 89–91; 2001, 106–8; Tetlock 2017)

5.3 Comparing the testimonies of multiple experts in personal finance

In Chapter 3.3, I analysed the heuristic value of 'counting agreement' of putative experts on a testimony, utilizing the fact that laypeople can compare different agents within a field. If multiple experts agree on a proposition and if laypeople have tentative reasons that each of them is trustworthy enough, then, so the intuition says, each independently agreeing party adds additional weight to the trustworthiness of the testimony from the perspective of the layperson. (see A. I. Goldman 2001, 97–104)

Regarding belief-independence, I defended Goldman's received view against recent criticism by David Coady, Lackey and Miller by arguing that his principle is being misunderstood as more demanding than it is. The simple principle holds that if a testimony is a perfect copy of another testimony, this additional testimony cannot justifiably increase our trust in the testimony. This should be seen as distinct from the possibility that copying testimony can also be seen as an affirmative meta-expert assessment. (see D. Coady 2006, 68–75; A. I. Goldman 2001, 97–104; Lackey 2013, 257–60; Miller 2013, 1311–14; Stegenga 2009)

Using agreements or even consensus of experts as an indicator of the trustworthiness of a proposition has been widely regarded as a useful heuristic within the philosophical debate, though it has not been undisputed. Although a potential for misguided inferences remains, the agreement still has epistemic value for laypeople next to other heuristics. To attain a more nuanced understanding of the value of expert agreement, I discussed accounts of its conditions by Miller and Goldman. I demonstrated that these approaches use a robustness-style argument and have therefore connected this debate with Stegenga's analysis thereof. The latter is an epistemic argument that says if the same proposition is being supported by several different sources, methods or reasonings then it is more robust and less likely driven by a mistake or bias (see Dietrich and Spiekermann 2023, chaps 2, 6.2; A. I. Goldman 2001, 97–104; Intemann 2017; Miller 2013, 1300–1314; Solomon 2007; Stegenga 2009)

Importantly, Stegenga asserts that such evidence needs to be both concordant, i.e. reasonably comparable with the other involved kinds of evidence, and independent, i.e. sensibly different and causally separated from the compared methods. There is a tension that the more independent these kinds of evidence are the more they tend to be discordant and vice versa., though these factors do not need to cancel each other out. Having this in mind, I argued that in the context of the expert-identification question we should understand robustness as addressing the relation of *testifiers* and not their specific methods since these would not be accessible to laypeople anyway. Accordingly, concordance is indicated by anything that makes the agreement of a putative expert more likely to be relevant for the question of interest, i.e. the agreeing expert is of the right discipline indicated by her educational background. Independent testimony, moreover, is indicated by anything that makes it more likely that an expert has reached their opinion by an independent process of reasoning, i.e. the expert has different methodological or theoretical commitments or has a different social identity background. (see Miller 2013, 1300–1314; A. I. Goldman 2001, 97–104; Stegenga 2009)

5.3.1 Assessing agreement between personal-finance experts

Having the discussion of Chapter 3.3 in mind, I now want to analyse whether it can be applied to a personal-finance context or whether there are differences we should be aware of. Beginning with what prevails, nothing in Chapter 4 contradicts the general epistemic intuitions of robustness-style arguments. The problem of finding independent evidential sources which should be concordant at the same time remains, but just as it did not render robustness impossible earlier neither does it now. Consequently, if a (hopefully representative) sample of personal-finance experts rather agrees on one testimony over another, the trustworthiness of this testimony is, *ceteris paribus*, more robust than the other. While the principal stands, I want to argue that, in detail, the peculiarities of personal finance very well affect the assessment of agreements between testifiers. (see A. I. Goldman 2001, 104–5; Stegenga 2009)

Since personal finance is informed by social sciences like economics and psychology, it is affected by a substantially higher intrinsic uncertainty for reasons I thoroughly discussed in Chapter 4. I distinguished that, while for basic expert tasks, this uncertainty is moderately high, it is immense for forecasting practices. Relevantly, higher intrinsic uncertainty causes more disagreements between (groups of) experts, since either an irregularity in a researched phenomenon elicits different perspectives or imperfections in the research method provoke different standpoints. Such disagreements would not be dishonest though laypeople would not be able to distinguish those from disagreements arising out of detrimental biases. Honest disagreements born out of intrinsic uncertainty are likely to constitute different reasonable research vectors. This closes the gap to a central observation from philosophers of science that disagreements and methods in social sciences are substantially more manifold than in the natural ones, though out of the social sciences economics is the most methodologically homogenous one. To continue, I want to ask whether this higher diversity in methods could promote robustness judgments in heuristic lay assessments due to the availability of more independent perspectives. (see Boumans and Martini 2014; Grable and Chatterjee 2022b; Hausman 2021, chap. 2.2-2.3; Lane 2014, 105–11; Martini 2014b, 123–28; Reiss 2014)

Interestingly, David Coady contends the opposite:

“Although it seems plausible that genuine experts in the natural sciences, or pure mathematics have a significant degree of independence from one another in their belief forming practices, experts on questions of history, and perhaps more generally, questions in the social sciences, are inevitably highly dependent on their sources.” (D. Coady 2006, 76)

Although he does not elaborate on this intuition any further, I think he focuses on the way scholars build upon the current state of their art. He might be concerned with the way scientists can do experiments to test hypotheses and add data to the debate. In this sense, he might perceive empirical studies in the natural sciences as less theory-laden than in the social sciences or that the choice for a sensible theoretical background is clearer in natural sciences. If this is the right interpretation, I assert that the diversity in social sciences is not the result of independence from theory or background assumptions but from there being multiple available pathways. Moreover, I think it would be premature to simply assume that empirical research in natural sciences is less theory-laden or less dependent on previous work. Furthermore, in historical research old testimonies are the research *object* and their investigation should not be conflated with the copying of another's research. (see D. Coady 2006, 75–76)

If we, then, accept that the many methodologies of social sciences increase the availability of more independent testimonial sources to compare for laypeople, does this come at the cost of concordance? After all, if these different methods are too idiosyncratic, then testimonies could not be properly compared and the gain in robustness would be meagre. Looking closer at how social sciences differ in methods, I want to argue that their diversity increases independence substantially more than they decrease concordance. Investigating (mainstream) economics and by implication personal finance, different methods often share the same core tools such as statistical measurements, norms of data collection, and ways of calculation. Also, Evans observed that economists mostly share a professional and methodological culture and rather disagree about the weight of variables and their interpretation. Accordingly, such experts can communicate their approaches and results with colleagues from other subdisciplines. Differences on the other hand rather occur in different theoretical bodies, framings of the researched phenomenon and background assumptions. Consequently, I argue that agreement between testifiers of different theoretical backgrounds is more independent than it involves discordance. In this vein, the more different backgrounds agree on the same proposition, *ceteris paribus*, the more robust testimonies can become. Would this, conversely, entail that higher intrinsic uncertainties improve the epistemic environment of a discipline, even if just from a lay perspective? After all, agreement is not an epistemic criterion for scientific experts but a fallible heuristic for laypeople. I do not think this is implied, instead robustness-style arguments partially compensate the heuristic value of a lack of a proper consensus. (see Evans 2014, 238–39; Martini 2014b; Stegenga 2009)

Moreover, laypeople need to be able to distinguish testifiers with different sub-disciplinary backgrounds. While we again should not expect laypeople to assess the content of these sub-disciplines, they can use indicators as discussed in Chapter 3.3.3. Without understanding, for example, what distinguishes a neoclassical economist from a monetarist, a layperson can learn that these are different theoretical frameworks, while still being relevant to the question at hand. In the context of personal finance, robustness would be indicated by experts seemingly addressing personal finance issues while arguing from different theoretical frameworks. (see Evans 2014, 238–45; Martini 2014b; Stegenga 2009)

Especially, personal-finance experts, however, can also agree or disagree due to non-epistemic reasons like (shared) biases. Since I argued in Chapter 5.1 that personal finance is closely entangled with the non-epistemic interests of experts and laypeople, laypeople should expect more detrimental interests present in this communication. Whether laypeople can profit from a robustness-style assessment of agreement between testifiers, then, especially depends on their ability to effectively discount biased testimonies. Besides the heuristics previously discussed in Chapter 5.1, this embellishes the relevance of Miller’s social diversity conditions. He argued that if putative experts with diverse social backgrounds agree on the same testimony, this, *ceteris paribus*, makes this testimony more robust against detrimental incentives. I would extend the notion of social backgrounds to any identities and apparent interests relevant to the issue. If their bias is too high to have epistemic value their excluded testimony at least does not dilute the robustness analysis. (Miller 2013, 1298–1300, 1311–14; Reiss and Sprenger 2020, chaps 2–3; Reiss 2014, 139–41, 149–51; Stegenga 2009)

Martini and Evans both promote the idea that robustness, due to the involvement of many different but still relevant putative experts, is a good substitute for consensus assessments in disciplines where no such clear consensuses are likely to manifest. I agree that in personal finance, amongst others, laypeople should search for robust agreements on positions instead of consensuses. Moreover, Philip Tetlock and Dan Gardener report the following: (see Evans 2014; Martini 2014b)

One also would have taken a hit for dismissing teams on the grounds that teams inevitably reduce independent perspectives and are inherently prone to groupthink. For it proved possible, especially among superforecasters, to engineer team norms in ways that promoted open-minded exchanges of ideas and out-performed people working alone.” (see Tetlock 2017, xxvi–xxvii; Tetlock and Gardner 2015)

This echoes my previous argument that independence and concordance do not need to cancel each other out. Tetlock and Gardener further add that aggregation algorithms provide more accurate predictions when they contain the beliefs of forecasters with diverse perspectives. The relevance of robustness in Tetlock's forecasting with very high intrinsic uncertainty further highlights the role of robustness in such contexts. Contrary to authors like Intemann, Solomon or Miller, I have offered an argumentation of how agreement or robustness can be of gradual value. These scholars focused on the fact that many disagreements, but rarely stable consensuses present in social sciences and how we should epistemically evaluate this aspect. My response was to show how testimonial agreement can be heuristically valuable without needing a proper consensus on a position by connecting Goldman's Bayesian intuition with the rationale of robustness. Hence, I fear scholars on consensuses in science forgot to entertain the idea that testimonial agreement can be of gradual heuristic epistemic value which is especially relevant in personal-finance contexts. (see A. I. Goldman 2001, 97–104; Intemann 2017; Miller 2013; Solomon 2007; Tetlock 2017, xxvii; Tetlock and Gardner 2015)

5.3.2 Strategically utilizing competition between personal-finance experts

Finally, I will consider how laypeople could use strategic heuristics to utilize multiple experts in personal finance. In recent chapters, I primarily used economic cheap talk models to analyse the strategic relationship between expert sender and lay receiver. Accordingly, I connected the insights from the single sender case to the philosophical debate to better understand how misaligned interests between experts and laypeople affect the former's trustworthiness. These models formally showed how too large misalignments turn truthful information transmission irrational, but moderate biases allow partial informativeness. I interpreted this cheap talk interaction as the expert's attempt to give misleadingly precise testimonies and laypeople's rational reaction to only believe the more general gist of the testimony, which improves the epistemic situation of laypeople at least somewhat. (see Crawford and Sobel 1982; Farrell and Rabin 1996)

In Chapter 3.3.4, I introduced a cheap talk model by Ambrus and Lu, which modified the original by contrasting one receiver with multiple senders. They contrived a strategy for the receiver, where she would check whether the two testimonies, she receives, are identical or not. If not, she knows that at least one has not been truthful since they share the same knowledge. Ambrus and Lu stipulate, that the receiver could choose an extreme reaction none of the experts could want, thereby threatening them into coordinating on the truth. I argued, however, that this

strategy constitutes an incredible threat, as the receiver would benefit more from realising a partially truthful equilibrium than from choosing a bizarre action that hurts both herself and the senders. This raises doubts on the means laypeople have in punishing experts and on the realistic applicability of this model. (see Ambrus and Lu 2014; Crawford and Sobel 1982; Sobel 2013)

While this issue is troubling, Ambrus and Lu have used this model to investigate scenarios where experts are not perfectly informed, i.e. the signals they learn involve probabilistic noise. This is interesting because in Chapter 4 I demonstrated that basic personal-finance practices involve higher levels of intrinsic uncertainty whereas economic forecasting even very high levels. Moreover, personal finance testimonies typically involve a wider extent of potential ambiguity which translates into larger state spaces and, by implication, larger hypothetical message and action spaces. Ergo, Ambrus and Lu's model extension could be especially informative for analysing the trustworthiness of personal-finance experts. (see Ambrus and Lu 2014; Crawford and Sobel 1982; Lane 2014, 105–11; Sobel 2013)

To look more closely into the rationale behind that point, I want to refer to Lane's study on uncertainty in (scientific) expert testimony. As I argued in Chapter 2.2, laypeople, unfortunately, cannot distinguish whether disagreements are caused by intrinsic uncertainty in the discipline or due to detrimental incentives, be they financial, reputational, or idealistic. Ergo, disagreement might not only indicate inevitable uncertainty for laypeople. In cases of honest testimony transmission, increased intrinsic uncertainty would already be expressed in more transparently ambiguous testimony, i.e. experts expressing less confidence or precision respectively. In the case of biased testimony, disagreements might arise out of intrinsic uncertainty 'or' detrimental incentives. This issue is one reason why intrinsic uncertainty may impede the laypeople's possibility to induce more honest testimony by comparing multiple sources.⁴⁵ (see Ambrus and Lu 2014; Lane 2014, 105–11; Sobel 2013, 317–21)

In this vein, Ambrus and Lu write:

“(...) equilibrium constructions provided in the literature require the senders to exactly reveal the true state, and punish individual deviations by rendering an action that is bad for both senders in the case of nonmatching reports. The latter is possible because after an out-of-equilibrium profile of messages by the experts, nothing restricts the beliefs of the receiver. Such constructions obviously break down if the experts observe the state with noise, however small, since nonmatching reports then occur along the equilibrium path.” (Ambrus and Lu 2014, 175)

⁴⁵ Lane cites empirical evidence that experts and laypeople both are systemically overconfident in judgments involving uncertainty, which entails the possibility that testimony in expertises with high uncertainty is unintentionally over-precise. Moreover, laypeople might tend to underestimate the (effects of) uncertainty in an expertise. (see Lane 2014, 105–11; Morgan and Mellon 2011)

Ambrus and Lu model the intrinsic uncertainty of expert knowledge by introducing noise to the senders' signals. This means that the expert information is not perfect and might contain errors to some degree, while the truth about the world is still the basis for the preferences of both the senders and the receiver. With a high positive probability strictly within 1 and 0, they observe the true signal or otherwise a random, state-independent variable.⁴⁶ If the noise parameter is small enough, the previous outcome of the noiseless game remains an equilibrium. After all, the expected utilities of the original strategy profiles remain similar enough to induce the same messages, threats, and actions. In simple terms, this means the proportionally closer to the truth the actual signals are, the more the outcome approximates a fully revealing equilibrium and its associated utility. This entails that if the noise is small enough, the testimony is as precise as the noise allows, leaving the testimony effectively unbiased though marginally ambiguous. (see Ambrus and Lu 2014, 178–81; Sobel 2013, 317–21)

Concerns regarding the rationality assumptions of this model aside suggest that while some intrinsic uncertainty is tolerable, too high amounts disable any chance to incentivise honesty by comparing the testimony of multiple senders. Since, substantial intrinsic uncertainty is not rare within sciences and other expertises, this punishing heuristic might not be widely applicable.⁴⁷ Consequently, this model endeavour formally supports doubts on the ability of laypeople to 'threaten' personal-finance experts into testifying honestly, since they cannot differentiate between disagreements due to intrinsic uncertainty or non-epistemic reasons. A point I have made repeatedly. While the long state space promotes the bargaining power of laypeople, intrinsic uncertainty in basic personal-finance expertise again reduces the possibility to induce truthfulness. It, however, is completely void in forecasting contexts since in scenarios of too high intrinsic uncertainty no fully revealing equilibrium is possible. (see Ambrus and Lu 2014, 178–81; Sobel 2013, 317–21)

There are, however, other ways multiple testifiers could improve the epistemic position of the layperson, namely by inducing competition in a market of personal-finance products and

46 Ambrus and Lu also analyse a second way to model intrinsic uncertainty by assuming signals for each state to depend on a continuous probability distribution. I will not consider this variety in my thesis. (see Ambrus and Lu 2014, 181–83)

47 To learn how much intrinsic uncertainty a discipline involves, laypeople could utilize the judgment of trusted meta-experts to distinguish different sources of disagreement. Laypeople could also observe additional (meta-)testimonies of putative experts on whether the disagreements of the debate arise out of intrinsic uncertainty or non-epistemic interests. (see C. A. J. Coady 1994, 203–23; Collins and Evans 2009, 60–79; A. I. Goldman 2001, 97)

biased advice. In Chapter 5.1.2, I introduced a model from Inderst and Ottaviani in which they analysed the incentive structure of commission-taking in personal-finance advice. In the same paper, they also discussed a variation of their model with multiple product provider–adviser dyads, i.e. two ‘senders’ and one receiver. In the process, they show an alternative way of how competition between experts can improve the epistemic situation of laypeople by focusing on how these dyads compete in a market with limited demand. (see Inderst and Ottaviani 2012, 405–7)

Starting from their baseline model without agency frictions (the nonnegativity constraint on the absolute compensation) or endogenized information collection, they add this second dyad, which competes with the first in selling an advanced product. In a market of wary customers, each dyad’s interest in increasing her market share decreases prices and, in turn, increases the surplus of the customer. This leads to a reduction of deadweight loss similar to Cournot duopolies while the testimony remains unbiased. More relevantly, market competition also affects the case of naïve clients. Even though testimony remains biased since the client’s ignorance makes it always profitable to urge them into buying advanced products, the dyads still need to compete on this biased price. Accordingly, the reservation value increases which reduces the scope firms have to bias advice. Firstly, the space between the naively anticipated product recommendation threshold (type threshold) and the true threshold is sinking, which indirectly necessitates closer-to-truth advice. Secondly, the generally feasible price range sinks which also caps the practical potential for exploitation. (see Inderst and Ottaviani 2012, 405–7)

Inderst and Ottaviani discuss two further aspects of competition with naïve clients. The dyads’ profits are still strictly higher when customers are naïve even though higher profits increase the elasticity of demand and thereby competition. After all, firms miss more utility from a naïve client than from a wary one and expanding the market share is more costly. Consequently, the customer’s naïve beliefs also dampen competition. The latter notwithstanding, the important point is that competition indirectly improves the epistemic situation of laypeople by capping the prices they could be otherwise compelled into and thus testifying closer to the truth. This is especially relevant since even careful laypeople are unlikely to perfectly anticipate detrimental incentives, as I have argued in Chapter 5.1.2. (see Inderst and Ottaviani 2012, 405–7)

Conclusion

In this thesis, I asked how laypeople could assess the trustworthiness of personal-finance advice along with the explicit consideration of strategic, game-theoretic thinking. In doing so, I assumed that laypeople are realistically limited in their resources to study the discipline or assess the experts and therefore need to rely on heuristic judgments. I characterised the latter as inferences to the best explanation, i.e. fallible reasons for laypeople that render best-possible belief more likely than chance. These ‘fast and frugal’ judgments depend on indicators perceived by laypeople that make it more likely that testimony is competent, honest or both. While expert knowledge is standardly inaccessible to laypeople, some parts of the content may be exoteric, i.e. witnessing the realisation of expert predictions as Goldman proposed. More importantly, accompanying features may be social facts such as monetary incentives, group affiliations or (dis-)agreements between experts who are accessible to laypeople, following Collins, Evans and Weinel’s sociology of expertise. (see Collins and Weinel 2011; Collins and Evans 2009; A. I. Goldman 2001; Miller 2013, 1299–1300; Wheeler 2020, chap. 7)

In a similar vein as Goldman, I argued that to meet the true interests of laypeople who want to take advantage of expert knowledge we should conceptualise experts as epistemic authorities. Accordingly, I have assumed experts to have real (cognitive) competencies substantially exceeding those of laypeople yet without any absolute minimum threshold, since laypeople could only profit from such agents. Moreover, I argued that experts relevant to my research question also fulfil a social function since they testify to laypeople. This becomes especially clear in my application to personal-finance advice, where I analyse lay assessments of experts who consult laypeople on their specific financial needs. Drawing on John Grable and Swarn Chatterjee, personal finance expertise is, in its modern form, a relatively young discipline informed by economics, psychology and sociology. I further categorized this advice into two groups of practices. Firstly, basic personal finance advice like saving strategies, debt consulting, tax optimisation and so on. Secondly, practices like investment advice depend on economic forecasting. (see A. I. Goldman 2018; Grable and Chatterjee 2022b)

In the same order, I started by investigating the effects of the social scientific background on basic testimonies of this field. According to central insights of the philosophy of economics and other social sciences, testimony is comparatively more causally complex, contingent, and uncertain compared to natural sciences like physics. Moreover, methods to explain social phenomena cannot be fully a-perspectival or free of norms of rationality. This normative entanglement and causal complexity increase the potential ambiguity of testimony along with

higher levels of intrinsic uncertainty. Finally, I suggested that the intimate relevance of these topics to both laypeople and experts further increases the chance for biases and detrimental incentives. Subsequently, I turned to the practice of economic forecasting and investigated its effectivity drawing on Tetlock's empirical studies of expert predictions and Angner's review of the overconfidence of economists. These showed that predictions of social phenomena quickly diminish in accuracy as the time horizon moves past two years until they are not better than chance. Especially, theorists championing one overarching theory seem to be systematically overconfident in their forecasts. (see Angner 2006; Reiss 2014; Tetlock 2017)

The first kind of heuristics I analysed, examines the honesty of experts by assessing their incentive structure. Here, I complemented the passive philosophical debate with game theoretic cheap talk modelling after Crawford and Sobel. This economic modelling formally supported the intuition that laypeople cannot expect truthful communication while their interests are misaligned with the expert. Importantly, however, if the bias is small enough, testimony can still be partially informative and lends itself to a gradual reception by laypeople. I further interpreted the strategic equilibrium as the attempt of experts to give misleadingly precise testimony while laypeople rationally ought to only believe the more general gist of the advice. Guided by Almassi's conditions for detrimental interests, I argued that biases should be expected to be more misleading when the potential ambiguity of testimony is high and experts have practical opportunities to act on their interests, i.e. there is a lack of institutional constraints on their dishonesty. Accordingly, how easy it is for putative experts to bias their testimony depends on the ability of laypeople to detect and punish an expert's misconduct. (see Almassi 2017; Crawford and Sobel 1982)

In the context of personal finance, I argued that the causal complexity of this discipline renders typical testimonies more ambiguous since more reasonable framings of one problem are possible, which increases the potential opportunities for experts to arbitrarily slant their message. Additionally, the combination of causal complexity and normative aspects in theory choice elicits increased intrinsic uncertainty in research methods. The latter impedes the detection and punishment of biases since laypeople cannot distinguish between disagreements due to intrinsic uncertainty or misaligned interests. Testimonies relying on forecasting, further, are affected by very high levels of intrinsic uncertainty which amplifies this problem. While laypeople cannot take advantage of the trust-inducing scientific institutions, which are often assumed in the general expert-identification debate, laypeople should cursorily investigate consumer protection laws and meta-expert control institutions in their jurisdiction. Nowadays,

for example, consumer protection in the EU is more extensive though it rather protects from fraud than warrants good advice. Since I am arguing that laypeople are unlikely to fully detect and disincentivise expert biases on their own, they have good reasons to avoid independent advice as good as possible. Inderst and Ottaviani showed that when laypeople underestimate expert biases, they are more likely to accept lower advising fees for more biased advice, which subsequently makes them worse off. Alternatively, I argued that lay people might strategically test experts by imagining how if their testimony were true and other people would know it, would this affect the effectiveness of the advice. If this hypothetical is not already the case, laypeople should question *why not*. (see Grundmann 2024; Inderst and Ottaviani 2012) (see Directive 2014/65/EU)

Moving further, I discussed Goldman's heuristic of observing 'salient successes' of expert competence to infer said competence. While this vector is indeed often helpful for laypeople, it is very contextual on the prevalence of such exoteric markers, like predictions laypeople can witness being realised or machines ostensibly working. Also, Goldman is aware, that not all such supposed indicators can allow an easy inference of expertise. Analysing the context of personal-finance advice, I argued that this heuristic is not at all applicable. The complex interplay of dynamically overlaying factors causing the performance of financial products is difficult to assign to specific expert interventions, especially for laypeople. Considering the meagre performance of financial forecasters' predictions, investment successes are most likely to be the result of either luck or pretence. (see Angner 2006; A. I. Goldman 2001, 106–8; Tetlock 2017)

Finally, I analysed heuristics that compare the testimony of multiple experts. While I concur that the overvaluing of agreement as a sign of trustworthiness could engender the danger of conflating social agreement with the likelihood of truth, I nevertheless concur with most scholars of the expert-identification debate that agreement can be of heuristic value to laypeople. Furthermore, I argued that accounts that promote this intuition like that of Goldman or Miller use robustness-style arguments. Drawing from Stegenga, I described the relevance of both independence and concordance of evidential sources. Even though both factors tend to contravene, evidence can be independent without being equally discordant. By observing social cues, laypeople can assess agreement towards a testimony regarding these two factors. Accordingly, concordance is indicated by anything that makes the agreement of a putative expert more likely to be relevant to the question of interest. Independent testimony, instead, is indicated by anything that makes it more likely that an expert has reached their opinion by an

independent process of reasoning. (see A. I. Goldman 2001, 97–104; Miller 2013; Stegenga 2009)

Agreement is also a valuable heuristic concerning putative *personal-finance* experts. The increased intrinsic uncertainty, however, makes it more difficult for laypeople to know whether (dis-)agreement is the result of limits in expertise or due to corrupting incentives. While strong consensuses are hardly possible here, robustness-style arguments can be compensatory heuristics for trustworthiness. The high diversity of, *prima facie*, reasonable theoretical background assumptions and frameworks allows the comparison of many independent perspectives. As long as this selection is, combined with other heuristic judgments, relevant enough for testimony, and testifiers are somewhat socially diverse, robustness can give laypeople reasons to sensibly assess the relative agreement of personal-finance experts towards a testimony and thus its trustworthiness. Since the intrinsic uncertainty in forecasting practices is especially high, this intensifies said results in that context. (see Angner 2006; Evans 2014, 139–45; A. I. Goldman 2001, 97–104; Miller 2013; Reiss 2014, 139–41; Stegenga 2009; Tetlock 2017)

The inability of laypeople to distinguish disagreement out of intrinsic uncertainty and non-epistemic interests is also problematic when analysing strategic heuristics. Accordingly, I investigated a heuristic of analysing disagreement between multiple testifiers in a modified cheap talk model by Ambrus and Lu. The authors showed formal support that when experts are perfectly informed, laypeople can infer from conflicting testimony that at least one testimony is not truthful and accordingly threaten both experts into communicating the truth by choosing an extreme option that would hurt everyone otherwise. I nevertheless argued that they use an incredible threat as the receiver in this context would prefer a partially truthful testimony over an extreme action. Regarding personal finance, Ambrus and Lu showed that laypeople might still benefit from this strategic heuristic, yet this is impossible in the highly uncertain contexts of economic forecasting. If we add financial product providers to the model like Inderst and Ottaviani, however, competition between multiple experts and providers might reduce the potential of exploiting the information asymmetry by ordinary mechanisms of market competition. (see Ambrus and Lu 2014; Inderst and Ottaviani 2012)

I hope to have sparked interest in the applied epistemology of assessing personal-finance experts and perhaps the endeavour to question the applicability of the general expert-identification debate to more particular domains of cognitive or even non-cognitive expertises. After all, I have argued that a sensible conceptualisation of ‘expertise’ notwithstanding,

different domains can have peculiarities that shape what is possible and important for laypeople in assessing trustworthiness. In this vein, future research could proceed in bridging gaps between general heuristic methods and other particular expert contexts. Certainly, the applied epistemology of personal finance or further economic experts could itself be deepened. For instance, future philosophical endeavours could specifically assess heuristics of rhetoric and demeanour of putative experts. Even though the general heuristic rationales might be the same as in most other (sales) contexts, analysing the specific ways they present themselves here, could help write guidelines for laypeople, which render them more alert to typical tactics of bad actors.

Moreover, I hope to have demonstrated the fruitfulness of employing perspectives of multiple disciplines on one problem area. Apart from the sociology of expertise (after Collins and Evans) and the philosophy of personal finance and economics, I especially engaged with economic communication models to supplement the philosophical debate with strategic, game-theoretic thinking. Whereas the philosophical debate addressed the relationship of experts and laypeople as static, the use of heuristics themselves can change the behaviour of experts which in turn can change what is the best action for laypeople. In this vein, interdisciplinary studies could highlight problems and solutions previous perspectives have not seen. (see Collins and Evans 2009)

Finally, I think the expert-identification debate could profit from a closer investigation of knowledge intermediaries and the role of the testimony chains. I deem this perspective to still be underdeveloped in this specific context. How such intermediaries learn from experts themselves, filter the information, and communicate them to laypeople seems to me to be of high relevance for discussing heuristics for laypeople.

Bibliography

‘About FPSB | FPSB’. n.d. Financial Planning Standards Board. Accessed 29 June 2024. <https://fpsb.org/about/>.

Almassi, Ben. 2017. ‘Toxic Funding? Conflicts of Interest and Their Epistemological Significance’. *Journal of Applied Philosophy* 34 (2): 206–20.

Ambrus, Attila, and Shih En Lu. 2014. ‘Almost Fully Revealing Cheap Talk with Imperfectly Informed Senders’. *Games and Economic Behavior* 88 (November): 174–89.

Anderson, Elizabeth. 2011. ‘Democracy, Public Policy, and Lay Assessments of Scientific Testimony’. *Episteme* 8 (2): 144–64.

Angner, Erik. 2006. ‘Economists as Experts: Overconfidence in Theory and Practice’. *Journal of Economic Methodology* 13 (1): 1–24.

Beek, Robert Van. 2022. ‘Personal Finance: A Practice Perspective’. In *De Gruyter Handbook of Personal Finance*, edited by John E. Grable and Swarn Chatterjee, 35–48. Berlin/Boston: De Gruyter.

Bekker, Immanuel. 1842. *Sextus Empiricus*. Reprint 2022. Berlin/Boston: Walter de Gruyter GmbH.

Berlin, Isaiah. 1953. *The Hedgehog and the Fox: An Essay on Tolstoy’s View of History*. UK: Weidenfeld & Nicolson.

Boumans, Marcel, and Carlo Martini. 2014. ‘Introduction: Experts and Consensus in Social Science’. In *Experts and Consensus in Social Science*, edited by Carlo Martini and Marcel Boumans, 1–13. Cham: Springer International Publishing.

Brewer, Scott. 1997. ‘Scientific Expert Testimony and Intellectual Due Process’. *Yale Law Journal* 107 (6): 1535–1682.

Bruin, Boudewijn de. 2021. ‘Epistemic Injustice in Finance’. *Topoi* 40 (4): 755–63.

Bruin, Boudewijn de, Lisa Herzog, Martin O’Neill, and Joakim Sandberg. 2023. ‘Philosophy of Money and Finance’. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta and Uri Nodelman, Spring 2023. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2023/entries/money-finance/>.

Campbell, John Y. 2006. ‘Household Finance’. *The Journal of Finance* 61 (4): 1553–1604.

Chakraborty, Archishman, and Rick Harbaugh. 2010. ‘Persuasion by Cheap Talk’. *American Economic Review* 100 (5): 2361–82. <https://doi.org/10.1257/aer.100.5.2361>.

Chase, James, and David Coady. 2019. ‘The Return of Applied Epistemology’. In *The Routledge Handbook of Applied Epistemology*, edited by David Coady and James Chase. New York: Routledge.

Coady, C.A.J. 1994. *Testimony: A Philosophical Study*. Online Edition. Oxford: Oxford University Press.

Coady, David. 2006. ‘When Experts Disagree’. *Episteme* 3 (1–2): 68–79.

———. 2012. *What to Believe Now: Applying Epistemology to Contemporary Issues*. 1. ed. New York: John Wiley & Sons, Ann Arbor.

- Collins, Harry. 1992. *Changing Order: Replication and Induction in Scientific Practice: With a New Afterword*. 1. [print.]. Chicago, Ill.: Univ. of Chicago Pr.
- Collins, Harry, and Robert Evans. 2009. *Rethinking Expertise*. Pbk. ed. 2011. Chicago, Ill.: Univ. of Chicago Press.
- Collins, Harry, and Martin Weinel. 2011. 'Transmuted Expertise: How Technical Non-Experts Can Assess Experts and Expertise'. *Argumentation* 25 (3): 401–13.
- Crawford, Vincent P., and Joel Sobel. 1982. 'Strategic Information Transmission'. *Econometrica* 50 (6): 1431–51.
- Croce, Michel. 2019. 'On What It Takes to Be an Expert'. *The Philosophical Quarterly* 69 (274): 1–21.
- Davidson, Roei. 2012. 'The Emergence of Popular Personal Finance Magazines and the Risk Shift in American Society'. *Media Culture & Society* 34 (1): 3–20.
- Dietrich, Franz, and Kai Spiekermann. 2023. 'Jury Theorems'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta and Uri Nodelman, Spring 2023. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2023/entries/jury-theorems/>.
- Duijf, Hein. 2021. 'Should One Trust Experts?' *Synthese* 199 (3): 9289–9312.
- Dulleck, Uwe, and Rudolf Kerschbamer. 2006. 'On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods'. *Journal of Economic Literature* 44 (1): 5–42.
- Elliott, Kevin C. 2014. 'Financial Conflicts of Interest and Criteria for Research Credibility'. *Erkenntnis* 79 (5): 917–37. <https://doi.org/10.1007/s10670-013-9536-2>.
- 'European Supervisory Architecture'. n.d. FMA Österreich. Accessed 29 June 2024. <https://www.fma.gv.at/en/international-affairs/european-supervisory-architecture/>.
- Evans, Robert. 2014. 'Expert Advisers: Why Economic Forecasters Can Be Useful Even When They Are Wrong'. In *Experts and Consensus in Social Science*, edited by Carlo Martini and Marcel Boumans, 50:233–52. Ethical Economy. Cham: Springer International Publishing.
- Farrell, Joseph, and Matthew Rabin. 1996. 'Cheap Talk'. *Journal of Economic Perspectives* 10 (3): 103–18.
- Fidler, Fiona, and John Wilcox. 2021. 'Reproducibility of Scientific Results'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2021. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/sum2021/entries/scientific-reproducibility/>.
- 'Financial Market Supervision in Austria'. n.d. FMA Österreich. Accessed 29 June 2024. <https://www.fma.gv.at/en/financial-market-supervision-in-austria/>.
- 'FMA Austria'. n.d. Front page. FMA Österreich. Accessed 29 June 2024. <https://www.fma.gv.at/en/>.
- 'FPSB'. n.d. Front page. Financial Planning Standards Board. Accessed 29 June 2024. <https://fpsb.org/>.
- Fricker, Miranda. 2007. *Epistemic Injustice*. Oxford University Press.

- Friedberg, Mark, Bernard Saffran, Tammy J. Stinson, Wendy Nelson, and Charles L. Bennett. 1999. 'Evaluation of Conflict of Interest in Economic Analyses of New Drugs Used in Oncology'. *JAMA* 282 (15): 1453–57.
- Fuller, Steve. 2018. 'Forecasting'. In *The Routledge Handbook of Applied Epistemology*, edited by David Coady and James Chase, 247–58. London: Routledge.
- Gibbons, Robert. 1992. *A Primer in Game Theory*. Essex: Pearson Education Limited.
- Goldberg, Sanford. 2011. 'If That Were True I Would Have Heard about It by Now'. In *Social Epistemology: Essential Readings*, edited by Alvin Goldman and Dennis Whitcomb, 92–108. Oxford: Oxford University Press.
- Goldman, Alvin I. 2001. 'Experts: Which Ones Should You Trust?' *Philosophy and Phenomenological Research* 63 (1): 85–110.
- . 2018. 'Expertise'. *Topoi* 37 (1): 3–10.
- . 2021. 'How Can You Spot the Experts? An Essay in Social Epistemology'. *Royal Institute of Philosophy Supplements* 89 (May): 85–98.
- Goldman, Alvin, and Cailin O'Connor. 2021. 'Social Epistemology'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2021. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2021/entries/epistemology-social/>.
- Goodwin, Jean. 2011. 'Accounting for the Appeal to the Authority of Experts'. *Argumentation* 25 (3): 285–96.
- Grable, John E., and Swarn Chatterjee, eds. 2022a. *De Gruyter Handbook of Personal Finance*. Berlin/Boston: De Gruyter.
- . 2022b. 'Defining Personal Finance'. In *De Gruyter Handbook of Personal Finance*, 1–16.
- Grasswick, Heidi. 2018. 'Feminist Social Epistemology'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2018. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2018/entries/feminist-social-epistemology/>.
- Grundmann, Thomas. 2024. 'Experts: What Are They and How Can Laypeople Identify Them?' In *Oxford Handbook of Social Epistemology*, edited by Jennifer Lackey and Aidan McGlynn, unpublished manuscript. Oxford: Oxford University Press.
- Hardwig, John. 1985. 'Epistemic Dependence'. *The Journal of Philosophy* 82 (7): 335–49.
- Hausman, Daniel M. 2021. 'Philosophy of Economics'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2021. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2021/entries/economics/>.
- Hogarth, Robin M., Tomás Lejarraga, and Emre Soyer. 2015. 'The Two Settings of Kind and Wicked Learning Environments'. *Current Directions in Psychological Science* 24 (5): 379–85.
- Hong, Lu, and Scott E. Page. 2004. 'Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers'. *Proceedings of the National Academy of Sciences* 101 (46): 16385–89.
- Hume, David. 2007. *An Enquiry Concerning Human Understanding*. Oxford World's Classics. Oxford: Oxford University Press.

- Inderst, Roman, and Marco Ottaviani. 2012. 'How (Not) to Pay for Advice: A Framework for Consumer Financial Protection'. *Journal of Financial Economics* 105 (2): 393–411.
- Intemann, Kristen. 2017. 'Who Needs Consensus Anyway? Addressing Manufactured Doubt and Increasing Public Trust in Climate Science.' *Public Affairs Quarterly* 31 (3): 189–208.
- Kahneman, Daniel. 2012. *Thinking, Fast and Slow*. Penguin Books: Penguin Psychology. London: Penguin Books.
- 'Karriere'. n.d. Verband Financial Planners. Accessed 29 June 2024. <https://www.cfp.at/afp/afp.nsf/sypages/karriere.html>.
- King, Mervyn, and John Kay. 2020. *Radical Uncertainty: Decision-Making for an Unknowable Future*. Hachette UK.
- Kitcher, Philip. 1995. *The Advancement of Science*. Paperback. New York: Oxford University Press.
- Kuhn, Thomas S. 2012. *The Structure of Scientific Revolutions*. 4. ed., 50th anniversary ed. 2015. Chicago, Ill.: Univ. of Chicago Press.
- LaBarge, Scoff. 1997. 'Socrates and the Recognition of Experts'. *Apeiron* 30 (4): 51–62.
- Lackey, Jennifer Amy. 2013. 'Disagreement and Belief Dependence: Why Numbers Matter'. In *The Epistemology of Disagreement*, edited by David Christensen and Jennifer Lackey, 243–68. Oxford: Oxford University Press.
- Lane, Melissa. 2014. 'When the Experts Are Uncertain: Scientific Knowledge and the Ethics of Democratic Judgment'. *Episteme* 11 (1): 97–118.
- Locke, John. 1979. *An Essay Concerning Human Understanding*. Clarendon ed., 1. issued in paperback. 2011. Oxford: Clarendon Press.
- Loewenstein, George, Cass R. Sunstein, and Russell Golman. 2014. 'Disclosure: Psychology Changes Everything'. In *Annual Review of Economics*, edited by K. J. Arrow and T. F. Bresnahan, 6:391–419. Palo Alto: Annual Reviews.
- Longino, Helen. 2019. 'The Social Dimensions of Scientific Knowledge'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2019. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/sum2019/entries/scientific-knowledge-social/>.
- Marshall, Alfred. 2009. *Principles of Economics: Unabridged Eighth Edition*. 8th ed. New York: Cosimo, Inc.
- Martini, Carlo. 2014a. 'Experts in Science: A View from the Trenches'. *Synthese* 191 (1): 3–15.
- . 2014b. 'Seeking Consensus in the Social Sciences'. In *Experts and Consensus in Social Science*, edited by Carlo Martini and Marcel Boumans, 115–30. Cambridge: Springer International Publishing.
- . 2020. 'The Epistemology of Expertise'. In *The Routledge Handbook of Social Epistemology*, edited by Miranda Fricker, Peter Graham, David Henderson, and Pedersen, Nikolaj J.L.L., 115–22. New York, NY: Routledge.
- Martini, Carlo, and Marcel Boumans, eds. 2014. *Experts and Consensus in Social Science*. Vol. 50. Ethical Economy. Cham: Springer International Publishing.

- Milgrom, P., and J. Roberts. 1986. 'Relying on the Information of Interested Parties'. *Rand Journal of Economics* 17 (1): 18–32.
- Milgrom, Paul R. 1981. 'Good News and Bad News: Representation Theorems and Applications'. *The Bell Journal of Economics* 12 (2): 380–91.
- Miller, Boaz. 2013. 'When Is Consensus Knowledge Based? Distinguishing Shared Knowledge from Mere Agreement'. *Synthese* 190 (7): 1293–1316.
- . 2020. 'The Social Epistemology of Consensus and Dissent'. In *The Routledge Handbook of Social Epistemology*, edited by Miranda Fricker, Peter Graham, David Henderson, and Nikolaj J.L.L. Pedersen, 230–39. New York: Routledge.
- Morgan, M. Granger, and Carnegie Mellon. 2011. 'Certainty, Uncertainty, and Climate Change'. *Climatic Change* 108 (4): 707.
- Oreskes, Naomi. 2012. *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*. Pbk ed. London: Bloomsbury.
- Pérez-González, Saúl, and María Jiménez-Buedo. 2023. 'Non Experts: Which Ones Would Trust You?' *Social Epistemology* 37 (5): 610–25.
- Quast, C. 2018. 'Towards a Balanced Account of Expertise'. *Social Epistemology* 32 (6): 397–418.
- Reiss, Julian. 2014. 'Struggling Over the Soul of Economics: Objectivity Versus Expertise'. In *Experts and Consensus in Social Science*, edited by Carlo Martini and Marcel Boumans, 131–52. Cham: Springer International Publishing.
- Reiss, Julian, and Jan Sprenger. 2020. 'Scientific Objectivity'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2020. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2020/entries/scientific-objectivity/>.
- Resnik, David B., and Kevin C. Elliott. 2013. 'Taking Financial Relationships into Account When Assessing Research'. *Accountability in Research* 20 (3): 184–205.
- Schickore, Jutta. 2022. 'Scientific Discovery'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta and Uri Nodelman, Winter 2022. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2022/entries/scientific-discovery/>.
- Schleiermacher, Friedrich. 1857. *Plato*. 3rd ed. Platons Werke. Berlin/Boston: Walter de Gruyter GmbH.
- Scholz, Oliver R. 2009. 'Experts. What They Are and How We Recognize Them. A Discussion of Alvin Goldman's Views'. *Grazer Philosophische Studien* 79 (1): 187–205.
- . 2018. 'Symptoms of Expertise: Knowledge, Understanding and Other Cognitive Goods'. *Topoi* 37 (1): 29–37.
- Schuchardt, Jane, Dorothy C Bagwell, William C Bailey, Sharon A DeVaney, and et al. 2007. 'Personal Finance: An Interdisciplinary Profession'. *Journal of Financial Counseling and Planning* 18 (1): 61–84.
- Sobel, Joel. 2013. 'Giving and Receiving Advice'. In *Advances in Economics and Econometrics: Tenth World Congress: Volume I: Economic Theory*, edited by Daron Acemoglu, Eddie Dekel, and Manuel Arellano, 1:305–41. Econometric Society Monographs. Cambridge: Cambridge University Press.

- Solomon, Miriam. 2007. *Social Empiricism*. Cambridge, Mass.; London: MIT Press.
- . 2015. *Making Medical Knowledge*. Oxford, New York: Oxford University Press.
- Spence, M. 1973. 'Job Market Signaling'. *Quarterly Journal of Economics* 87 (3): 355–74.
- Stegenga, Jacob. 2009. 'Robustness, Discordance, and Relevance'. *Philosophy of Science* 76 (5): 650–61.
- Stichter, Matt. 2015. 'Philosophical and Psychological Accounts of Expertise and Experts'. *HUMANA.MENTE Journal of Philosophical Studies* 8 (28): 105–28.
- Tadelis, Steven. 2013. *Game Theory: An Introduction*. Princeton University Press.
- 'Tasks of the FMA'. n.d. FMA Österreich. Accessed 29 June 2024. <https://www.fma.gv.at/en/tasks-of-the-fma/>.
- Tetlock, Philip E. 2017. 'Expert Political Judgment: How Good Is It? How Can We Know? - New Edition'. In *Expert Political Judgment*, 2nd ed. Princeton; Oxford: Princeton University Press.
- Tetlock, Philip E., and Dan Gardner. 2015. *Superforecasting: The Art and Science of Prediction*. Random House.
- 'Verband'. n.d. Verband Financial Planners. Accessed 29 June 2024. <https://www.cfp.at/afp/afp.nsf/sysPages/verband.html>.
- 'Verband Financial Planners'. n.d. Front page. Verband Financial Planners. Accessed 29 June 2024. <https://www.cfp.at/>.
- Walton, Douglas Neil. 1997. *Appeal to Expert Opinion: Arguments From Authority*. University Park, PA, USA: Pennsylvania State University Press.
- Watson, Jamie Carlin. 2020. *Expertise: A Philosophical Introduction*. New York: Bloomsbury Publishing.
- Weber, Max. 2011. *Methodology of Social Sciences*. New York: Routledge.
- Wheeler, Gregory. 2020. 'Bounded Rationality'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2020. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2020/entries/bounded-rationality/>.
- Wittgenstein, Ludwig. 1998. *Philosophical Investigations*. Blackwell.
- Zagzebski, Linda Trinkaus. 2012. *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief*. Oxford University Press.

Appendix

English Abstract: How could laypeople assess the trustworthiness of personal-finance advice?
An interdisciplinary study.

Following Alvin Goldman's seminal paper (2001), philosophers are debating again how laypeople could be rationally justified to trust experts. This is a riddle in applied epistemology as laypeople lack sufficient time, energy, and education to check the expert's competence or learn every relevant expertise themselves, but still would profit from utilizing the best available knowledge. Philosophers have, so far, discussed several fallible heuristic strategies laypeople could use to gauge trustworthiness. I argue that the debate has implicitly focused on expertises in natural sciences and therefore, I will investigate whether popular heuristics work equally well in social-science-informed, service expertises like personal finance, which is of intimate relevance to many people. Moreover, I suggest that the philosophical debate neglected strategic effects in the expert-lay relationship, which I will supplement with perspectives from economic game theory. Focusing on heuristics of 'detecting and punishing detrimental incentive structures', 'identifying indicators of salient success', and 'assessing the agreement of multiple experts to a testimony', my interdisciplinary study will show that the heightened ambiguity, intrinsic uncertainty, and proximity to detrimental monetary interests in personal finance creates a difficult epistemic environment for laypeople. While laypeople cannot identify any salient success indicators here, they should focus on assessing incentive structures by anticipating the strategic effects of commission-taking, consumer protection laws and control institutions. Within the putative experts with the best incentive structure, they should trust the testimony with the most, relevant agreement.

German Abstract: How could laypeople assess the trustworthiness of personal-finance advice?
An interdisciplinary study.

Seit Alvin Goldmans einflussreichem Journalartikel (2001) debattieren Philosophen erneut darüber, wie Laien Experten *rational* vertrauen könnten. Das ist ein Rätsel der angewandten Epistemologie, da Laien nicht genügend Zeit, Energie und Bildung haben, um die Kompetenzen von Experten zu überprüfen oder sich selbst jegliches relevante Fachwissen anzueignen. Jedoch würden sie vom richtigen Expertenwissen profitieren. Diese Philosophen haben bisher mehrere fehlbare, heuristische Strategien diskutiert, die Laien verwenden könnten, um mutmaßliche Experten einzuschätzen. Ich argumentiere, dass sich die Debatte implizit auf Expertenwissen

in den Naturwissenschaften konzentriert hat und werde daher untersuchen, ob populäre Heuristiken bei sozialwissenschaftlich fundierten, service-orientierten Expertisen, wie Personal Finance, ebenso gut funktionieren würden. Außerdem kritisiere ich, dass die philosophische Debatte die strategischen Effekte der Experten-Laien-Beziehung vernachlässigt, weswegen sie mit ökonomischer Spieltheorie ergänzen werde, um diese Perspektiven zu ergänzen. Dabei fokussiere ich mich auf Heuristiken wie, das ‚Erkennen und Bestrafen problematischer Interessensstrukturen‘, das ‚Identifizieren von offensichtlichen Erfolgsindikatoren‘ und das ‚Abwägen von Zustimmungen mehrerer Experten zu einer Frage‘. Meine interdisziplinäre Studie wird zeigen, dass die erhöhte Mehrdeutigkeit, die intrinsische Unsicherheit und die Nähe zu problematischen monetären Interessen, Personal Finance Expertise zu einem, schwierigen epistemischen Umfeld für Laien macht. Laien können in diesem Kontext keine Erfolgsindikatoren identifizieren und sollten sich daher auf die Bewertung von Interessensstrukturen konzentrieren, indem sie die strategischen Effekte von Provisionsabgaben, Verbraucherschutzgesetzen und Kontrollinstitutionen antizipieren. Unter den vermeintlichen Experten mit den besten Anreizen sollten Laien dann die Meinung mit der größten, relevanten Zustimmung vertrauen.