

Evidence of Evidence as Higher-Order Evidence

Anna-Maria A. Eder and Peter Brössel

Contents

1	Introduction	2
2	Evidence of Evidence as First-Order Evidence	3
2.1	Fitelson’s Counterexamples to the EEE Slogan	4
2.2	Tal and Comesaña’s Specifications of the EEE Slogan	5
3	Evidence of Evidence as Higher-Order Evidence	8
3.1	Dyadic Bayesianism, Higher-Order Evidence, and Two Assumptions	8
3.2	Evidence of Evidence as Higher-Order Evidence	11
4	New Specifications of the EEE Slogan	15
4.1	Evidence of More Comprehensive Evidence	15
4.2	Evidence of Complementary Evidence	17
5	Conclusion	21

Abstract

In everyday life and in science we acquire evidence of evidence and based on this new evidence we often change our epistemic states. An assumption underlying such practice is that the following **EEE Slogan** is correct: ‘evidence of evidence is evidence’ (Feldman 2007, p. 208). We suggest that evidence of evidence is best understood as *higher-order evidence* about the epistemic state of agents. In order to model evidence of evidence we introduce a new powerful framework for modelling epistemic states, **Dyadic Bayesianism**. Based on this framework, we then discuss characterizations of evidence of evidence and argue for one of them. Finally, we show that whether the **EEE Slogan** holds, depends on the specific kind of evidence of evidence.

Keywords Evidence; Evidence of Evidence; Higher-Order Evidence; Evidential Support; Social Epistemology

1 Introduction

Human agents are social beings who exchange information constantly. The beliefs and credences that agents communicate to us influence our own doxastic states. This is so in cases in which we share the same body of evidence with them, and also in cases in which we don't. Sometimes we learn from others and sometimes we even defer to the doxastic attitudes of agents for whom we have evidence that they have evidence that we lack. When our partner is looking in the fridge searching for milk, and returns empty-handed saying that there is no milk in the fridge, we assume that there is no milk in the fridge and that our partner has evidence for this. When a physician is looking at an X-ray of our arm, we believe that the physician has evidence for whether our arm is broken and we follow her recommendation for treatment. In other cases we suspend judgement, for instance when other agents have evidence that we don't have and which counter-supports our beliefs or credences. Feldman provides an example of such a latter case:

Criminal Case Example Consider [...] the example involving the two suspects in a criminal case, Lefty and Righty. Suppose now that there are two detectives investigating the case, one who has the evidence about Lefty and one who has the evidence incriminating Righty. They each justifiably believe in their man's guilt. And then each finds out that the other detective has evidence incriminating the other suspect. If things are on a par, then suspension of judgment is called for. (Feldman 2007, p. 208)

An assumption underlying the practices in these examples is that the following slogan, coined by Feldman, is in some way or other correct:

EEE Slogan '[E]vidence of evidence is evidence' (Feldman 2007, p. 208).

Feldman is aware that this is an imprecise slogan and immediately offers a more precise specification according to which 'evidence that there is evidence for p is evidence for p ' (Feldman 2007, p. 208; notation adjusted). Despite its intuitive appeal, there have been many critical discussions of the **EEE Slogan** (e.g. Comesaña and Tal 2015, Fitelson 2012, Moretti 2016, Roche 2014). Typically these discussions propose particular specifications of *evidence of evidence* and the **EEE Slogan**, and then show that the slogan—given a particular specification of *evidence of evidence*—is incorrect. Most, if not all, counterexamples brought forward against such specifications exploit a feature of evidential support that is fairly well known: evidential support—understood as increase in probability—is not transitive. Attempts to specify the **EEE Slogan** concentrate on identifying conditions under which evidential support is transitive (e.g. Moretti 2015, Roche 2014, Tal and Comesaña 2017).

The aim of the present paper is to provide a nuanced approach towards the **EEE Slogan**. In pursuing this aim we do not focus on general conditions under which evidential support is transitive. Our focus is narrower: first and foremost, we provide a new account of what evidence of evidence is. Then in a second step we provide an account of when and why we are permitted or even ought to change our credence when we have evidence that some other agent possesses evidence for a proposition. We provide an account of when and why the **EEE Slogan** holds.

We proceed as follows. In section 2 we introduce and criticize discussions of the **EEE Slogan**. Most critical discussions of the slogan portray evidence of evidence as if it were additional first-order evidence and it is argued that the **EEE Slogan** would require that the evidential support relation is transitive, which then is shown not to be transitive. We conclude the section by suggesting that evidence of evidence is more adequately understood as *higher-order evidence* about the epistemic states of agents, and that it should be treated differently from first-order evidence. In section 3, we introduce **Dyadic Bayesianism**, which is a new powerful framework for modelling (rational) epistemic states of agents. According to the framework, an agent's (rational) epistemic state is best modelled not by the agent's credences alone, which would be standard in Bayesian epistemology. It is best modelled by a dyad consisting of (i) the agent's (rational and probabilistic) reasoning commitments, which reflect the justificatory import of the evidence and capture how the agent justifies her credences, and (ii) her total evidence. Based on **Dyadic Bayesianism**, we discuss characterizations of evidence of evidence that refer to it as higher-order evidence, and we argue for one of them. Building on our preferred characterization of evidence of evidence, we develop a new approach towards the **EEE Slogan** in section 4. We investigate when and how evidence about epistemic states of agents should influence our own doxastic attitudes when we have evidence that they possess evidence for some proposition. We conclude by reviewing our results in section 5.

2 Evidence of Evidence as First-Order Evidence

Most intuitive cases in which we would like to refer to the **EEE Slogan** are formally reconstructed as if they involved three first-order pieces of evidence. The **EEE Slogan** is then often understood as in some way or another presupposing the transitivity of evidential support. For instance, (a variation of) our milk example from the beginning would be reconstructed as follows:

Milk Example Hanna acquires evidence e_1 , the proposition that her partner Carina has studied the content of the fridge and claims that there is no milk in the fridge. Evidence e_1 then supports the proposition p_1 , the proposition that it seems to Carina that there is no milk in the fridge. The proposition p_1 then supports the proposition p_2 , the proposition that there is no milk in the fridge.

Most philosophers would acknowledge that p_1 differs in kind from the other propositions because it is about the epistemic state of an agent and not about the world. Nevertheless, for understanding the **EEE Slogan** many think the differences can be ignored or idealized away. Ignoring any difference in kind between p_1 and the other propositions, they then would argue as follows: e_1 not only supports p_1 but also p_2 . The alleged reason is this: e_1 is evidence for p_1 , which is evidence for p_2 and, therefore and in line with the **EEE Slogan**, e_1 is also evidence for p_2 . The **EEE Slogan** seems to require transitivity of evidential support.

It is standard in the literature on the **EEE Slogan** to understand evidential support as incremental support. Following the literature on the slogan we adopt such an understanding of evidential

support as well.¹ *Incremental support* is in terms of increase in probability: a proposition p is (incrementally) supported by a proposition e just in case the probability of p given e is above the probability of p alone (i.e. $\Pr(p|e) > \Pr(p)$). Unfortunately, understood as incremental support, evidential support is not transitive in general. Fitelson’s (2012) counterexamples to specifications of the **EEE Slogan** demonstrate this indubitably. In the following we discuss the simplest of his counterexamples to illustrate the basic idea behind such examples.

2.1 Fitelson’s Counterexamples to the EEE Slogan

Fitelson (2012) asks us to consider the following example:²

Card Example Let p_1 be the proposition that c is a black card, p_2 the proposition that c is the ace of spades, and p_3 the proposition that c is an ace. The proposition p_1 is evidence for p_2 and p_2 is evidence for p_3 . They each increase the probability of the subsequent proposition. We immediately see that p_1 and p_3 are probabilistically independent (assuming a standard card deck). (Adapted from Fitelson 2012, p. 85)

This shows that evidential support, understood as incremental support, is intransitive. Fitelson now turns this into a counterexample against the **EEE Slogan**. He is aware that examples like the one above do not belong to the intended application of the **EEE Slogan**, since the proposition p_2 does not seem to have the right kind of content; it is not about evidence or what an agent believes or knows. He finds a simple and ingenious way to turn **Card Example** into a counterexample against the **EEE Slogan**. The following is a variation of his original counterexample:

Another Card Example Again, let p_1 be the proposition that c is a black card, p_2 the proposition that c is the ace of spades, and p_3 the proposition that c is an ace. Let us assume that it is part of our background knowledge that Manuel knows exactly what c is. This ensures that the following holds for all possible bodies of evidence e : $\Pr(p_2|e) = \Pr(\text{Manuel knows } p_2|e)$. Thus, p_1 supports that *Manuel knows* p_2 (since p_1 supports p_2), and that *Manuel knows* p_2 supports p_3 . It is still not true that p_1 supports p_3 . (Adapted from Fitelson 2012, pp. 85–86)

Formally, Fitelson’s counterexample against the **EEE Slogan** makes use of the fact that evidential support is not transitive: that p_1 supports *Manuel knows* p_2 and *Manuel knows* p_2 supports p_3 does not guarantee that p_1 supports p_3 . By suggesting the right kind of interpretation of these propositions, Fitelson turns this observation into a counterexample. In particular, formally the proposition *Manuel knows* p_2 is treated as if it were a first-order proposition p_2 with no special content. However, it is interpreted as a proposition about which evidence is available to an agent. Fitelson also discusses other specifications of the **EEE Slogan** that seem plausible but can be refuted by similar strategies and depend on the fact that, formally, evidence of evidence is treated

¹Most of our results do not depend on this understanding of evidential support. Indeed, often the more straightforward interpretation of our results is in terms of *absolute support*, which is in terms of high probability: a proposition p is (absolutely) supported by a proposition e just in case the probability of p given e is high, that is, above an appropriate threshold t , where $t > .5$ (i.e. $\Pr(p|e) \geq t_{>.5}$). The understanding of evidential support as incremental support and the understanding as absolute support are the two main understandings of evidential support in the literature.

²The presentation of his example is adapted here.

as first-order evidence and that evidential support is not transitive.

Since Fitelson published his counterexamples, epistemologists have focused on specifying conditions under which evidential support is transitive. This is motivated by the assumption that the **EEE Slogan** would be subject to the same conditions. For example, Roche (2014) suggests—based on results from Bayesian confirmation theory that trace back to Keynes 1921 and Shogenji 2003—that evidence of evidence is evidence under screening-off. He makes this suggestion based on the fact that evidential support is transitive given such a condition.³ Clearly, if evidential support were transitive, we could not provide counterexamples such as **Another Card Example**. Tal and Comesaña (2017) argue against Roche’s approach and suggest that the correct way of specifying the **EEE Slogan** is in terms of defeaters. Based on this idea, Tal and Comesaña provide one of the most precise and sophisticated accounts for hedging the **EEE Slogan**. Despite the sophisticated nature of their account, it too exemplifies what is going wrong in the present debate about the **EEE Slogan**: they too treat evidence of evidence as if it were first-order evidence. They do so at least in their formal reconstruction of cases involving evidence of evidence. In the following we study their account closely and learn from their mistakes.

2.2 Tal and Comesaña’s Specifications of the EEE Slogan

Tal and Comesaña start their approach with the following observation:

[W]e believe that the correct way of fixing EEE principles is in terms [...] of defeaters. Notice a very interesting thing: when e is evidence for p which is evidence for q but e is at the same time a defeater for the support that p provides for q , the positive screening-off condition will not be satisfied. That e is a defeater for the support that p provides for q means that $e \wedge p$ is not evidence for q , whereas the positive screening-off condition requires that the support that $e \wedge p$ gives to q not be lower than the support that p alone provides to q . If $e \wedge p$ provides no support at all to q , then of course that support will be lower than the one that p alone provides given that p provides any positive support, as is required in all EEEs. We have therefore found a partial explanation of the evidential relevance of the positive screening-off condition: it rules out cases where e is itself a defeater for the support that p provides to q . But given that the positive screening-off condition is overkill, we should replace it with a no-defeaters condition. (Tal and Comesaña 2017, p. 108)

The idea is this: if the agent’s evidence e supports the proposition p , that there is a proposition e^* that is true and that supports q , and if e is not a defeater for the support that p provides to q , then e itself supports q to some positive degree. Based on this idea, they propose the following specification of the **EEE Slogan**:

Tal and Comesaña’s EEE Suppose ‘ $S(p, q, +\alpha)$ ’ says that proposition p supports proposition q to some positive degree α and ‘ $T(p)$ ’ says that p is true, then the following holds:

³For literature on the transitivity of evidential support given screening-off conditions see, for instance, Keynes 1921, Roche 2014, and Shogenji 2003.

$$S(e, \exists e^*[T(e^*) \wedge S(e^*, p, +\alpha)], +\beta) \wedge S(e \wedge \exists e^*[T(e^*) \wedge S(e^*, p, +\alpha)], p, +\gamma) \rightarrow \exists \delta S(e, p, +\delta)$$

We shall argue that **Tal and Comesaña's EEE** is incorrect and that their understanding of evidence of evidence is not helpful. Let us discuss **Tal and Comesaña's EEE** by means of the following example:⁴

New Card Example Let p_1 be the proposition that c is a card in the suit of spades and p_2 be the proposition that card c is an ace. Now suppose there are two agents, Elisabeth and Andreas, who start with the same a priori credence function Pr . Then Elisabeth but not Andreas receives a piece of evidence e_1 saying that Vincent claims to know that c is an ace of spades, that is, Vincent claims to know $p_1 \wedge p_2$. Then, in a second step, Elisabeth and Andreas receive the proposition p_1 as evidence. Note that p_1 taken by itself is irrelevant for whether or not p_2 .

Before discussing the example in a formal way let us discuss it informally. First, intuitively, Elisabeth receives evidence of evidence for p_2 whereas Andreas does not. Second, intuitively, Elisabeth's total evidence (i.e. $e_1 \wedge p_1$) is evidence for p_2 , since it is evidence of evidence for p_2 (for this case we would like to employ the **EEE Slogan**); Andreas's total evidence (i.e. p_1) is irrelevant for p_2 (the **EEE Slogan** is not applicable here). The reasoning behind both statements is the following. Elisabeth's total evidence includes e_1 : that Vincent claims to know that $p_1 \wedge p_2$ and the evidence p_1 . These pieces of evidence together support that Vincent indeed has evidence for $p_1 \wedge p_2$ and that Vincent really knows $p_1 \wedge p_2$. The intuitions behind the **EEE Slogan** lead us to believe that Elisabeth possesses evidence of evidence for p_2 as well. Andreas only received evidence p_1 which is irrelevant for p_2 and the intuitions behind the **EEE Slogan** seem not to support that the obtained evidence is evidence of evidence.

Now let us discuss this example in more detail, in terms of the Bayesian theory of evidential support. (We assume that the agents involved are rational from a Bayesian perspective.) We can make two observations here: first, note that both agents already know a priori that there is a proposition e^* that support p_2 to some positive degree $+\alpha$, simply because it implies the latter proposition. In particular, propositions such as $p_1 \wedge p_2$ and also *Vincent knows* ($p_1 \wedge p_2$) imply p_2 and, thus, support p_2 to some positive degree $+\alpha$. Let us use the log-likelihood measure⁵ to measure evidential support, then $\alpha = \log\left[\frac{\text{Pr}(e|p_2)}{\text{Pr}(e|\neg p_2)}\right] = \infty$, if $e \models p_2$ (see Fitelson 1999). Thus, according to Elisabeth's and also Andreas's a priori credence function Pr : $\text{Pr}(S(p_1 \wedge p_2, p_2, \infty)) = 1$, $\text{Pr}(S(\textit{Vincent knows } (p_1 \wedge p_2), p_2, \infty)) = 1$ and $\text{Pr}(\exists e^*[S(e^*, p_2, \infty)]) = 1$. The only thing Elisabeth and Andreas do not know a priori is whether any of these propositions that support p_2 to degree ∞ are also true.

Second, note that the evidence that both Elisabeth and Andreas receive (i.e. p_1) provides evidential support for the truth of propositions that support p_2 to degree ∞ . In particular, p_1 (which we

⁴Our discussion takes up an argument by Moretti (2016), which is worked out in more detail here.

⁵The log-likelihood measure is one of the most prominent Bayesian confirmation measures in the literature. This is due to its flexibility, the possibility to factor out independent pieces of evidence, and its close relationship to the Bayes factor. See Brössel 2013 and especially Fitelson 1999, 2001 for a discussion of this measure.

assumed to be irrelevant for p_2), supports $p_1 \wedge p_2$ (and $T(p_1 \wedge p_2)$), which implies that p_2 . Thus, both agents receive evidence p_1 which supports the conjunction $T(p_1 \wedge p_2) \wedge S(p_1 \wedge p_2, p_2, \infty)$ and the corresponding existential statement $\exists e^*[T(e^*) \wedge S(e^*, p_2, \infty)]$. For both agents, the evidence p_1 also supports something else. In particular, p_1 supports that Vincent knows that $p_1 \wedge p_2$. In particular *Vincent knows* ($p_1 \wedge p_2$) implies p_2 and p_1 . Therefore, p_1 supports *Vincent knows* ($p_1 \wedge p_2$) and *Vincent knows* ($p_1 \wedge p_2$) supports p_2 . (Arguably, however, Elisabeth receives much stronger evidence in favour of *Vincent knows* ($p_1 \wedge p_2$), after all Elisabeth’s total evidence includes the evidence e_1 that Vincent claims to know $p_1 \wedge p_2$.)

In the following we present two problems that Tal and Comesaña have to face. The first problem for **Tal and Comesaña’s EEE** is this: intuitively, only Elisabeth receives evidence of evidence for p_2 and, thus, evidence for p_2 . Andreas’s evidence should not be understood as evidence of evidence and the evidence received by Andreas should be irrelevant for p_2 . According to **Tal and Comesaña’s EEE**, the opposite is true. Both agents receive evidence of evidence. Elisabeth receives the evidence $e_1 \wedge p_1$ and Andreas the evidence p_1 . Both pieces of evidence support the existential claim $\exists e^*[T(e^*) \wedge S(e^*, p_2, \infty)]$. Supporting such an existential claim is what distinguishes ordinary evidence from evidence of evidence, according to **Tal and Comesaña’s EEE**. This trivializes the characterization of evidence of evidence: all evidence is evidence of evidence for arbitrary p . This trivializes the notion of evidence of evidence for p .

The second problem for **Tal and Comesaña’s EEE** is that they can hardly explain why Elisabeth’s evidence $e_1 \wedge p_1$ supports p_2 , while Andreas’s evidence p_1 does not. We saw that both agents’ evidence supports the existential claim $\exists e^*[T(e^*) \wedge S(e^*, p_2, \infty)]$ (and they do so for the same propositions $p_1 \wedge p_2$ and *Vincent knows* ($p_1 \wedge p_2$)). Thus, they would need to argue that Andreas’s evidence is a defeater for the evidential support p_2 receives from the existential claim $\exists e^*[T(e^*) \wedge S(e^*, p_2, \infty)]$. However, this is impossible: that there is a true proposition that implies p_2 , implies that p_2 is true. Logical implication provides support that cannot be defeated (at least not in the present framework).⁶ Thus, according to Tal and Comesaña’s specification of the **EEE Slogan** we would expect that both agents receive evidence that supports p_2 —this is contrary to what we have been arguing.

Taken together, the above two problems show that **Tal and Comesaña’s EEE** is inadequate as a specification of our intuitive **EEE Slogan**. Although the discussion in Tal and Comesaña 2017 enriches the debate tremendously by bringing to the foreground propositions that speak about evidential support, we still have nothing to help us explain the intuitive appeal of the **EEE Slogan**. It seems to us that the problem here is that we have been operating with an inadequate characterization of evidence of evidence. According to the characterization with which we have been operating, evidence e is evidence of evidence for p if e supports that there exists a true proposition e^* which supports p . We believe the mistake is to allow the proposition e , the evidence of evidence for p , to be first-order evidence about the world, as long as e supports the existence of further true

⁶We do not think that the original evidence p_1 functions in this case as higher-order evidence, à la Christensen 2010, that calls our logical abilities into question, and thus would defeat the support p_2 receives from the existential claim.

propositions that support p . But we have just seen that every proposition e supports the existence of true propositions that support p .

Based on previous observations we propose a new characterization of evidence of evidence. Instead of characterizing evidence of evidence as first-order evidence and running the risk that one can show—yet again—that specifications of the **EEE Slogan** can be trivialized or shown to be inadequate, we propose to characterize evidence of evidence as higher-order evidence.

3 Evidence of Evidence as Higher-Order Evidence

In the previous section, we proposed to understand evidence of evidence as higher-order evidence. Now we need to present our detailed account of how to understand such higher-order evidence. In the present debate, we propose to understand evidence of evidence as higher-order evidence *about the epistemic state of an agent*. Thus, in a first step, we say more about how to model the epistemic states of agents. In a second step, we introduce candidates for characterizations of evidence of evidence.

3.1 Dyadic Bayesianism, Higher-Order Evidence, and Two Assumptions

Dyadic Bayesianism

In Brössel and Eder 2014 we argued that for social epistemological purposes it does not suffice to identify the epistemic state of an agent with the agent’s doxastic attitudes, like credences. To understand the dynamics of rational reasoning in social settings, we argued, we need a more fine-grained framework for modelling the epistemic states of agents. We proposed such a framework, which we here refer to as **Dyadic Bayesianism**.⁷ This framework is also apt for present purposes. We introduce it first informally, and then in a formal way.

According to **Dyadic Bayesianism** the following holds. First, the (rational) *epistemic state of an agent* is best modelled by a dyad consisting of (i) the agent’s (rational) reasoning, or confirmation,⁸ commitments—which reflect the justificatory import of the evidence and capture how the agent justifies her credences—and (ii) the agent’s total evidence. We refer to the latter component as *the evidential state of the agent*. Second, the agent’s (rational) credences equal her (rational) reasoning commitments conditional on the total evidence—which captures the idea that the agent uses her reasoning commitments to form her credences by conditionalizing on the evidence available to her. Finally, third, (rational) credences and (rational) reasoning commitments obey the probability calculus. This implies that the agents that we are considering are Bayesian agents, who have ideal epistemic states. (For the appropriate normative interpretation of these probabilities see Eder ms.) Brössel (accepted) shows that the framework of **Dyadic Bayesianism** is fruitful for providing a satisfactory account of confirmation. In Brössel and Eder 2014, we show that on the

⁷In Brössel and Eder 2014 we refer to it as *Pluralistic Bayesianism*.

⁸In Brössel and Eder 2014 we use the term *confirmational commitments*. However, now we think that it is more apt to use the term ‘reasoning commitments’.

basis of **Dyadic Bayesianism** we can provide a powerful account for finding agreement among agents that initially disagree with each other. We show that with the account one can satisfy many requirements which alternative accounts that deal with peer disagreement do not satisfy. In this paper we want to employ **Dyadic Bayesianism** to provide an account of evidence of evidence.

Putting it in a more precise and formal way, **Dyadic Bayesianism** says the following:

Dyadic Bayesianism An agent s 's (rational) epistemic state is

1. a dyad/ordered-pair $ES_s = \langle Pr_{R_s}, tev_s \rangle$ consisting of (i) s 's reasoning commitments, Pr_{R_s} , and (ii) s 's total evidence, tev_s ,
2. such that s 's credences are as follows: $Pr_{Cr_s}(p) = Pr_{R_s}(p|tev_s)$,
3. and both Pr_{Cr_s} and Pr_{R_s} obey the probability calculus.

Higher-Order Evidence

In this paper we understand evidence of evidence as higher-order evidence. We characterize higher-order evidence as follows:

Higher-Order Evidence A proposition e is higher-order evidence for some agent s if and only if there is some agent s^* such that e describes, or evaluates, or guides the (formation of the) epistemic state of s^* .

Note first that we assume that judgements of the (ir-)rationality of epistemic states are evaluative or guiding in the sense relevant here. Furthermore, a consequence of the conception of higher-order evidence employed here is that propositions stating logical relations or evidential support relations are not higher-order evidence (provided the latter is understood as being independent of the agent's reasoning commitments). Such propositions may become logically equivalent to higher-order evidence in our sense if we add certain presuppositions about the normativity or rationality of logic and evidential support. Without such presuppositions, however, they are not higher-order evidence because they neither describe, nor evaluate, nor guide an agent's epistemic state.⁹

Before we proceed and characterize evidence of evidence, we emphasize that it is one strength of **Dyadic Bayesianism** that we can distinguish different kinds of higher-order evidence that play an important role in epistemology. First, we can distinguish between higher-order evidence about evidence and higher-order evidence about reasoning commitments. The former describes, or evaluates, or guides the evidential state of an agent, the latter describes, or evaluates, or guides the reasoning commitments of the agent. And then there is higher-order evidence about evidence *and* reasoning commitments. In particular, higher-order evidence about the credences of an agent

⁹Our **Higher-Order Evidence** corresponds to a narrow conception of higher-order evidence. There might be a broader conception according to which, for instance, evidence about logical or evidential support relations counts as higher-order evidence, independently of any additional presuppositions. In any case, most, if not all, of what we argue for would hold if we adopted a wider conception of higher-order evidence or rephrased **Higher-Order Evidence** in such a way that the right side of the biconditional expresses only a sufficient condition for higher-order evidence. We prefer to work with this more informative biconditional.

is such evidence. In our framework, the credences of an agent are (by definition) formed by updating her reasoning commitments on her available evidence. Information about that agent's credences, thus, informs us about both her reasoning commitments and her evidence. Second, we can distinguish between intra- and interpersonal higher-order evidence. Intrapersonal higher-order evidence describes, or evaluates, or guides the epistemic state of oneself. Interpersonal higher-order evidence describes, or evaluates, or guides the epistemic state of another agent.

In individual epistemology, intrapersonal higher-order evidence about one's evidence plays an important role when we study the epistemic role of, for example, perception and memory. This kind of higher-order evidence is crucial for determining the input for our own reasoning. What it describes, or evaluates, or guides is not how we reason, but on what we (should or are permitted to) base our reasoning. Intrapersonal higher-order evidence about reasoning commitments in turn plays an important role when we want to study inferential reasoning. This kind of higher-order evidence describes, or evaluates, or guides how we reason, and thus is relevant for determining what credences or beliefs we (should or are permitted to) hold on the basis of our evidence.

Interpersonal higher-order evidence is crucial in social epistemology. In Brössel and Eder 2015 we discuss in detail how to revise one's credence in the light of interpersonal higher-order evidence that describes the reasoning commitments of another agent and displays disagreement with the agent. In this paper we place interpersonal higher-order evidence about the evidential state of another agent at the centre of our debate about evidence of evidence. We discuss this kind of higher-order evidence in detail in the following.

Two Assumptions

To isolate the question of how to deal with evidence of evidence from other issues in epistemology, we want to factor out other questions that are closely related. To this end we introduce two idealizing assumptions.

First, we add an assumption about the evidence available to the agents. Following other authors working on the **EEE Slogan** (e.g. Tal and Comesaña 2017), we assume that the total evidence of agents does not include false propositions. This assumption simplifies our task considerably. Among others, we can ignore questions concerning whether one is justified in relying on the evidence available to other agents or to oneself. In addition, it ensures that the evidential states of two agents can always be combined to obtain a more comprehensive evidential state that is consistent. Thus, the assumption allows us to set aside the epistemological problems of perception and testimony.¹⁰ An obvious consequence of the possibility of combination is that the question of how to deal with *specified* evidence of evidence can be answered outright: if an agent s receives the

¹⁰Perception is the primary source of information about the world. Testimony is the primary source of information about the evidence and beliefs of other agents. Given our assumptions we can ignore the question of whether and how strongly we are justified in accepting such perceptual and testimonial evidence. See, for example, Brössel 2017, Lyons 2016, and Pryor 2000 for a discussion of the epistemological problems of perception. See Adler 2017 and Lackey and Sosa 2006 for a discussion of the epistemological problems of testimony.

higher-order evidence e that the total evidence of another agent s^* , tev_{s^*} , contains the *specified* proposition e^* , then s should update her evidential state by including e^* in her total evidence. The remaining question is how to deal with evidence of *unspecified* evidence, namely evidence that states that some other agent has some evidence for p without specifying the exact evidential proposition that the other agent has. In the following we focus on evidence of unspecified evidence.

Second, we make the idealizing assumption that both agents—the first agent s , who receives higher-order evidence of the second agent s^* 's evidence, as well as s^* , who has first-order evidence in support of a proposition—have the same reasoning commitments (i.e. $\text{Pr}_{R_s} = \text{Pr}_{R_{s^*}}$) and that both of them are aware of this. For this reason, for the remainder of this paper we drop the subscripts for the agents when we speak of reasoning commitments. A consequence of this assumption and **Dyadic Bayesianism** is that both agents agree in their credences when they share the same body of total evidence. (Remember: according to **Dyadic Bayesianism** an agent's credence in a proposition equals the agent's reasoning commitments with respect to the proposition conditional on the total evidence of the agent.) This allows us to ignore questions concerning whether one is justified in relying on some other agent's reasoning capacities or rather on one's own reasoning capacities, and whether it is possible that agents have different but equally rational reasoning commitments. It also allows us to ignore questions concerning how to deal with cases of peer disagreement, namely cases in which equally rational agents who share the same body of evidence assign different credences to the same proposition in light of the same body of evidence—questions such as: Should one be steadfast or change one's credences in such a case? And if so, how should one revise one's credences? We want to avoid these debates here, but at the same time acknowledge that a comprehensive theory of evidence of evidence—which we cannot present in this paper—would engage with those questions. (We have already answered the latter question on how to change one's credence when one is required to find agreement in Brössel and Eder 2014.)

3.2 Evidence of Evidence as Higher-Order Evidence

Now we are in a position to discuss possible characterizations of evidence of evidence as higher-order evidence that render the **EEE Slogan** correct.¹¹

A First Attempt at Characterizing Evidence of Evidence

With the above account of epistemic states to hand, we are in a position to introduce characterizations of evidence of evidence for some proposition p . Before we concentrate on evidence of evidence for some proposition p , let us consider what makes evidence into evidence of evidence. We submit the following characterization:

¹¹That we focus on this kind of evidence does not exclude that we are open to the possibility of there being different adequate characterizations of evidence of evidence. However, in this paper we are interested in particular characterizations that make the **EEE Slogan** correct. For an alternative approach see Dorst and Fitelson ms.

A Characterization of EE A proposition e is (higher-order) evidence of evidence for some agent s if and only if there is some agent s^* with epistemic state $ES_{s^*} = \langle \text{Pr}_R, tev_{s^*} \rangle$ such that either

1. there is some (specified) proposition e^* such that e states that e^* is included in the total evidence of s^* (i.e. $tev_{s^*} \models e^*$) [evidence of specified evidence]
- or
2. e states that there is some (unspecified) proposition e^* that is included in the total evidence of s^* (i.e. $tev_{s^*} \models e^*$) [evidence of unspecified evidence].

Now we can turn to the question of when such evidence of evidence is evidence of evidence for some proposition p . The main idea is this: we obtain evidence e of evidence e^* for a proposition p just in case e says that some agent s^* possesses some piece of evidence e^* that supports p .

We are interested in when and why we are permitted or even ought to change our credence when we have evidence that some agent possesses a piece of evidence. Often, we receive such evidence about other agents and not about ourselves. Thus, we focus on evidence of evidence that is interpersonal higher-order evidence of evidence. Strictly speaking, evidence of evidence can be intrapersonal evidence of evidence too. For example, we might obtain higher-order evidence saying that we forgot an unspecified proposition e^* that at an earlier time was implied by our evidence and that supported and still supports p . Arguably, this is evidence of evidence too. Nevertheless, in this paper we restrict ourselves to interpersonal higher-order evidence about evidential states and ignore this possibility. (If one feels this restriction is too severe, then one can simply treat agents at different time points as if they were different agents.) As a first pass at the notion of evidence of evidence for p let us consider the following characterization.

A Characterization of EE for p A proposition e is (higher-order) evidence of evidence for a proposition p for some agent s with epistemic state $ES_s = \langle \text{Pr}_R, tev_s \rangle$ if and only if there is some (other)¹² agent s^* with epistemic state $ES_{s^*} = \langle \text{Pr}_R, tev_{s^*} \rangle$ such that

1. there is some (specified) proposition e^* such that e states that e^* is included in the total evidence of s^* (i.e. $tev_{s^*} \models e^*$) and that e^* supports p (i.e. $\text{Pr}_R(p|e^*) > \text{Pr}_R(p)$). [evidence of specified evidence for p]
- or
2. e states that there is some (unspecified) proposition e^* that is included in the total evidence of s^* (i.e. $tev_{s^*} \models e^*$) and that this e^* supports p (i.e. $\text{Pr}_R(p|e^*) > \text{Pr}_R(p)$). [evidence of unspecified evidence for p].

This seems to be an apt characterization of evidence of evidence for some proposition p if we want to characterize it in a very general manner. It is important to note, however, that **A Characterization of EE for p** certainly does not render the **EEE Slogan** correct. There are various reasons for this: some concern misleading evidence, others shared evidence, and some concern the fact that it is easy to acquire evidence that supports a proposition. We explain this in the following.

¹²We add ‘(other)’ to emphasize our focus on interpersonal higher-order evidence.

According to **A Characterization of EE for p**, evidence of evidence that is known to be misleading evidence for a proposition is still evidence for the proposition. However, (higher-order) evidence of evidence for p that is known to be misleading evidence for p is certainly not evidence for that proposition. Consider the following example:

Party Example Suppose you learn e_1 , which states that Ann's total evidence includes some unspecified evidence e^* that supports the proposition p_1 that Ben will attend your party. However, Ben just called you to cancel because his partner has surprised him with concert tickets. In this case you do not take e_1 to be evidence for p_1 . Instead you would take e_1 as what it is: higher-order evidence of Ann's misleading evidence for p_1 .

Evidence of evidence for a proposition that is known to be misleading is certainly not evidence for the proposition: so the **EEE Slogan** is incorrect, given **A Characterization of EE for p**.

There are also other cases that make clear that the **EEE Slogan** is incorrect when evidence of evidence for p is understood as suggested by **A Characterization of EE for p**. Consider the following case, which involves evidence of shared evidence:

Another Party Example Suppose your total evidence supports the proposition p_1 that Ben will attend your party. You then obtain higher-order evidence e_1 that Ann has evidence that Ben will attend the party and that Ann and you *share* exactly the same evidence (including this higher-order evidence).

You receive (higher-order) evidence of evidence for the proposition that Ben will attend the party. But this would not change your credences in the proposition. This higher-order evidence does not support the proposition and, intuitively, it should not increase your credence in the proposition. Hence, the **EEE Slogan** would be obviously incorrect when evidence of evidence for p is understood as suggested by **A Characterization of EE for p**.

There is another way to see that the **EEE Slogan** cannot be correct when evidence of evidence for p is understood as suggested by **A Characterization of EE for p**. This is so because it is (too) easy to obtain (higher-order) evidence of evidence for a proposition p . In particular, when we receive evidence of evidence, almost always we receive evidence of evidence for p . Suppose you receive evidence that another agent s^* 's total evidence tev_{s^*} implies a proposition e^* . Then you also obtain evidence that this agent's total evidence implies $(e^* \vee p)$, which almost always supports p . ($(e^* \vee p)$ is a logical consequence of p and logical consequences of some proposition always support it.) So almost always, evidence of evidence comes with some evidence of evidence for p , and we certainly do not want to say that all of this evidence is evidence for p , whatever p is. In addition, a characterization of evidence of the evidence of another agent should not be one where insignificant parts of this agent's entire body of total evidence are taken into consideration, it should be evidence of evidence that is determinant for the agent's current credence in that proposition. More specifically, given this evidence, the rest of the agent's total evidence should be (probabilistically) irrelevant for the agent's credence in the proposition in question.

Our Characterization of Evidence of Evidence

As emphasized several times, the entire investigation here is devoted to answering when and why we may or should revise our credences in the light of evidence of evidence. For this reason, we are interested in higher-order evidence of *new* or *unknown* evidence for p , that is to say evidence of evidence for p that is part of the total evidence of some (other) agent, but that is *unknown* to us. For this reason we are also interested in evidence of evidence that is sufficiently significant for our credence in p , and ideally sufficiently relevant to determine our credence in p . Only when the unspecified evidence is sufficiently significant for p for the agent who possesses it, can it be significant for the agent who merely receives higher-order evidence of the existence of such evidence for p .

In the following we rephrase **A Characterization of EE for p** in line with the above considerations. The new characterization of evidence of evidence understands evidence of evidence as (higher-order) evidence which states that there exists an *unknown* piece of evidence for a proposition that is possessed by another agent. (Remember that **A Characterization of EE for p** does not require that the evidence is unknown to the agent who receives the evidence of evidence for p .) The other agent's credence in the proposition given the agent's total evidence is higher than one's credence in it given one's total evidence. This reflects that the other agent's total evidence increases the support of the proposition in comparison to one's total evidence. In accordance with this, we propose to revise the above characterization as follows:

Our Characterization of EE for p A proposition e is (higher-order) evidence of *unknown* evidence for a proposition p for some agent s with epistemic state $ES_s = \langle \text{Pr}_R, \text{tev}_s \rangle$ if and only if there is some (other) agent s^* with epistemic state $ES_{s^*} = \langle \text{Pr}_R, \text{tev}_{s^*} \rangle$ such that

1. there is some (specified) proposition e^* such that e states that proposition e^* is included in the total evidence of s^* (i.e. $\text{tev}_{s^*} \models e^*$) but not in the total evidence of s (i.e. $\text{tev}_s \not\models e^*$) such that $\text{Pr}_{C_{r_{s^*}}}(p) = \text{Pr}_R(p|e^*) > \text{Pr}_R(p|\text{tev}_s) = \text{Pr}_{C_{r_s}}(p)$. [evidence of unknown specified evidence for p]
or
2. e states that there is some (unspecified) proposition e^* included in the total evidence of s^* (i.e. $\text{tev}_{s^*} \models e^*$) but not in the total evidence of s (i.e. $\text{tev}_s \not\models e^*$) such that $\text{Pr}_{C_{r_{s^*}}}(p) = \text{Pr}_R(p|e^*) > \text{Pr}_R(p|\text{tev}_s) = \text{Pr}_{C_{r_s}}(p)$. [evidence of unknown unspecified evidence for p].

With this characterization to hand, we can study which specification(s) of the **EEE Slogan** can be turned into a tenable principle. Before doing so, let us explain why **Our Characterization of EE for p** captures one important and intuitive conception of *evidence of evidence* correctly. First, not just any evidence about the evidential state of (another) agent is evidence of evidence in the above sense. Given the above characterization, it is required that the higher-order evidence e claims that there is some (specified or unspecified) evidence e^* included in the evidential state of s^* that is unknown to the agent s . Since this evidence is unknown to s it might change s 's credence in p . Second, the evidence (specified or unspecified) e^* is determinant for the agent s^* 's

credence, this being ensured by the assumption that $\Pr_{Cr_{s^*}}(p) = \Pr_R(p|e^*)$; given e^* , the rest of s^* 's total evidence is probabilistically irrelevant to p . **Our Characterization of EE for p** is a very strong conception of evidence of evidence for p . For many possible usages, **A Characterization of EE for p** might be more fruitful, especially if we want to distinguish various forms of evidence of evidence, independently of the question whether these notions underwrite a version of the **EEE Slogan**. However, **Our Characterization of EE for p** is a very promising one for finding a correct specification of the **EEE Slogan**.

4 New Specifications of the EEE Slogan

Given **Our Characterization of EE for p**, we know in advance that evidence of evidence for p informs the agent s that the total evidence of the second agent s^* contains some proposition not already contained in the evidence of s . This assumption amounts to: $tev_s \not\equiv tev_{s^*}$. Now we need to distinguish two cases. In the first case, the second agent s^* possesses all the evidence that the first agent s already possesses and more, that is: $tev_{s^*} \models tev_s$. In this case the agent s receives—what we refer to as—*evidence of more comprehensive evidence*. In the second case, s herself possesses evidence that the second agent s^* does not possess, that is: $tev_{s^*} \not\models tev_s$. In this second case s receives—what we refer to as—*evidence of complementary evidence*.

4.1 Evidence of More Comprehensive Evidence

Case 1 $tev_{s^*} \models tev_s$ (and $tev_s \not\models tev_{s^*}$)

Suppose agent s obtains the higher-order evidence e which states that s^* possesses all the evidence of agent s and some additional unspecified proposition e^* that is not included in the total evidence of s and that $\Pr_{Cr_{s^*}}(p) = \Pr_R(p|tev_{s^*}) > \Pr_R(p|tev_s) = \Pr_{Cr_s}(p)$. This is a case in which the second agent s^* is better informed than the first agent s and has a higher credence in p . s^* possesses more evidence and, after receiving the higher-order evidence e , s knows this. Should agent s increase her credence in p in the light of the higher-order evidence e ?

Without our idealizing assumption that both agents have identical reasoning commitments, we do not think that the answer should be an unequivocal ‘Yes’. If, for whatever reason, s mistrusts the reasoning commitments of agent s^* , then s should not increase her credence in p . Obtaining evidence of evidence that s^* takes to be relevant for p is not evidence for p if we have reasons to believe that s^* is mistaken about whether her more comprehensive evidence is probabilistically relevant for p . For example, very often patients have more evidence about their illnesses or injury than physicians; the patient knows where it hurts and how much, how badly her ankle bent when she landed, etc. But even then, physicians rarely feel they should defer in their judgement to the credences of the patient. Instead they convince patients to share their evidence so that they can apply their (medical experts) reasoning commitments to it (and if necessary commission further inquiries and tests). The reason for this is simple: physicians mistrust the reasoning commitments

of their patients, and more often than not they are right in doing so.

Given our simplifying assumption that both agents have identical reasoning commitments, we would answer the above question (i.e. should the agent s increase her credence in p in the light of the higher-order evidence e ?) with an unequivocal ‘Yes’. The reason is that if both agents employ the same reasoning commitments, then the **EEE Slogan** has a similar spirit as van Fraassen’s (1984) reflection principle,¹³ which we want to rephrase here in our terminology as a coherence principle for reasoning commitments.

Reflection for Reasoning Commitments An agent s ’s reasoning commitments Pr_R commit her to a credence r with respect to a proposition p conditional on the following evidence: there is a proposition e^* such that s ’s reasoning commitments concerning s ’s credence in p given e^* say that it should be r and that e^* is her future total evidence,

More formally: $\text{Pr}_R(p|\exists e^*[\text{Pr}_R(p|e^*) = r \wedge e^* = \text{tev}_s^{\text{Future}} \wedge e^* \models \text{tev}_s^{\text{Current}}]) = r$
if s is certain that her future self is not misinformed (i.e. has no false evidence).

This reflection principle is obviously relevant for the **EEE Slogan**: we only need to assume that it is irrelevant whose evidence e is—that is, whether it is agent s ’s future evidence or agent s^* ’s current evidence. In this paper we set aside questions concerning the reliability and trustworthiness of agents, and have assumed that all evidence is true. Given this assumption, it should be irrelevant whose evidence it is as long as it is logically stronger and, thus, more informative.¹⁴ In such cases, it seems advisable to defer to the other agent’s credences; they are better informed than we are and they are using the same reasoning commitments to form their credences. In the light of these assumptions, we can slightly reformulate the principle to obtain a tenable specification of the **EEE Slogan**:

Evidence of More Comprehensive Evidence is Evidence An agent s ’s reasoning commitments Pr_R commit her to a credence r with respect to a proposition p conditional on the following

¹³The standard formulation of this principle is in terms of credences and reads as follows:

Reflection Agent s ’s current credence in p on the condition that her future credence in p is r , should be r ,

Formally: $\text{Pr}_{C_{r_s}^{\text{Current}}}(p|\text{Pr}_{C_{r_s}^{\text{Future}}}(p) = r) = r$

if the following three conditions are satisfied:

1. s is certain that her future self is not misinformed (no false evidence).
2. s is certain that her future self has more information than she has now.
3. s believes that her future credence is obtained by rational updating.

In the literature this principle is often interpreted as a coherence principle between your current credences and your anticipated future credences (Huttegger 2013, Titelbaum forthcoming).

Our **Reflection for Reasoning Commitments** below relies on the idea that if the agent’s future credences are different, this must be due to the fact that the agent will receive some unspecified evidence and that she will change her credences in response to this evidence as required by **Dyadic Bayesianism**, namely by forming her credences by conditionalizing her reasoning commitments on her total evidence including this unspecified future evidence. The latter assumption is the reason why we do not need to assume that one’s future credences are obtained by rational updating. Instead, we can be explicit about the assumption that one’s future evidence implies one’s old evidence (and is thus more informative). Our **Reflection for Reasoning Commitments** also makes it obvious that this is a coherence principle for one’s reasoning commitments. It requires the commitment to reason in accordance with your reasoning commitments whatever unspecified evidence will come along in the future.

¹⁴Cartesians in the sense of Barnett (forthcoming) might uphold this or a similar principle even if give up the assumption that all evidence is true, since they defend perceptual impartiality. Dogmatists like Pryor (2000) presumably will ultimately reject it, because they believe we should give some kind of priority to our own perceptual experiences.

evidence: there exists a proposition e^* such that s 's reasoning commitments concerning s 's credence in p given e^* say that it should be r and that e^* is the more comprehensive total evidence of another agent s^* .

More formally: $\Pr_R(p|\exists s^*\exists e^*[\Pr_R(p|e^*) = r \wedge e^* = tev_{s^*} \wedge e^* \models tev_s^{Current}]) = r$

This specification of the **EEE Slogan** seems to us as plausible as the above **Reflection for Reasoning Commitments** (and van Fraassen's reflection principle discussed in n. 13), and it is subject to similar qualifications and restrictions. In addition, we see how evidence of more comprehensive evidence for p can provide you with relevant information to increase your credence in p . Suppose e is higher-order evidence of more comprehensive evidence for p . Then e claims that there is some (specified or unspecified) proposition e^* included in the total evidence of s^* (i.e. $tev_{s^*} \models e^*$) but not in the total evidence of s (i.e. $tev_s \not\models e^*$) and that $r = \Pr_{Cr_{s^*}}(p) = \Pr_R(p|tev_{s^*}) > \Pr_R(p|tev_s) = \Pr_{Cr_s}(p)$. If we now add our assumption that $tev_{s^*} \models tev_s$, e implies the proposition that we need to apply the above principle: $\exists s^*\exists e^*[\Pr_R(p|e^*) = r \wedge e^* = tev_{s^*} \wedge e^* \models tev_s^{Current}]$. And when we apply the principle we defer to the agent s^* with more evidence. Thus, higher-order evidence of (specified or unspecified) more comprehensive evidence for p is evidence for p .

4.2 Evidence of Complementary Evidence

Let us consider cases in which both agents possess evidence that the other agent does not possess; here both agents possess complementary evidence.

Case 2 $tev_{s^*} \not\models tev_s$ (and $tev_s \not\models tev_{s^*}$)

The majority of instances in which we would like to refer to the **EEE Slogan** to determine whether we should revise our credences are such that both agents have complementary pieces of evidence. We have already discussed some of these instances: first, recall Feldman's **Criminal Case Example**, where he asked us to consider 'two suspects in a criminal case, Lefty and Righty' and the 'two detectives investigating the case, one who has the evidence about Lefty and one who has the evidence incriminating Righty' (Feldman 2007, p. 208). Feldman then notes that, upon finding out that the other detective has evidence incriminating the other suspect, each detective is required to suspend judgement. Second, consider the example of Hanna observing Carina studying the content of the fridge. Presumably, Carina does not have all the evidence that Hanna, her partner, has. When Carina returns from the fridge she possesses an important piece of evidence that Hanna is missing. In the following we discuss whether one should revise one's credences in the light of (higher-order) evidence of unknown (specified or unspecified) complementary evidence. Before we do so, let us consider how a Bayesian agent forms her credences in the light of her total evidence.

Assume again that we measure evidential support with the help of the log-likelihood measure, l , of confirmation or evidential support (Fitelson 1999), which we introduced before. It is known that the agent's credence in p in the light of the agent's total evidence tev_s is a function from her reasoning commitments $\Pr_R(p)$ (which correspond to some extent to her a priori credence in p)

and the evidential support p receives from tev_s , $l(p, tev_s)$.¹⁵ The higher the support $l(p, tev_s)$, the higher the resulting credence.

Let us first study how we could incorporate evidence from other agents if we receive evidence of *unknown* specified evidence for p . (This will later also help us to better understand how to revise one's credence in the light of evidence of *unknown*, *unspecified*, complementary evidence.) In particular, let this evidence say that the specified proposition e^* is part of the total evidence of the second agent s^* and suppose agent s 's total evidence is tev_s . Remember, we assume that all evidence is true and that one can fully rely on the pieces of evidence included in someone's evidential state. Thus, agent s would simply add e^* to her current total body of evidence and update her credences accordingly. Then one can prove the following for the log-likelihood measure of evidential support (Brössel 2015):

Theorem

$$l(p, tev_s \wedge e^*) = l(p, tev_s) + l(p, e^*) + \log \left[\frac{\frac{\Pr_R(tev_s \wedge e^* | p)}{\Pr_R(tev_s | p) \times \Pr_R(e^* | p)}}{\frac{\Pr_R(tev_s \wedge e^* | \neg p)}{\Pr_R(tev_s | \neg p) \times \Pr_R(e^* | \neg p)}} \right]$$

This theorem shows that the effect of adding e^* to s 's total evidence concerning the proposition p depends on three factors:

1. to which degree tev_s supports p :

$$l(p, tev_s) = \log \left[\frac{\Pr_R(tev_s | p)}{\Pr_R(tev_s | \neg p)} \right]$$

2. to which degree e^* supports p :

$$l(p, e^*) = \log \left[\frac{\Pr_R(e^* | p)}{\Pr_R(e^* | \neg p)} \right]$$

3. to which degree the total evidence of s and the new evidence e^* cohere with each other in the light of p and in the light of $\neg p$ (in the sense of probabilistic coherence introduced in Shogenji 1999):

$$\log \left[\frac{\frac{\Pr_R(tev_s \wedge e^* | p)}{\Pr_R(tev_s | p) \times \Pr_R(e^* | p)}}{\frac{\Pr_R(tev_s \wedge e^* | \neg p)}{\Pr_R(tev_s | \neg p) \times \Pr_R(e^* | \neg p)}} \right]$$

If the sum of the latter two factors is greater than 0, then adding e to s 's total evidence would result in obtaining evidential support for p , and thus to increasing s 's credences in p (when compared with s 's credences in p given tev_s alone). If they are negative, then this would result in counter-support, and thus in a decrease of s 's credence (when compared with s 's credences in p

¹⁵The exact expression is not required in the main part of the paper, and would hinder readability:

$$\Pr_{C_{r_s}}(p) = \Pr_R(p | tev_s) = \frac{e^{\frac{1}{2}l(p, tev_s)} \Pr_R(p)}{e^{\frac{1}{2}l(p, tev_s)} \Pr_R(p) + e^{\frac{1}{2}l(\neg p, tev_s)} \Pr_R(\neg p)}$$

given tev_s alone).

This shows that even if we assume **Our Characterization of EE for p** and restrict the application of the **EEE Slogan** to evidence e of *unknown specified* evidence e^* for p , the **EEE Slogan** is not correct in general. If the proposition e^* does not cohere with the rest of the one's total evidence in the light of p , then this additional evidence will not increase our credence. And since the evidence e of evidence reports to the agent s which specified proposition e^* is included in the other agent s^* 's evidence, s herself needs to determine what her credence in p should be in the light of $e^* \wedge tev_s$.

Let us ask whether one should revise one's credence in p even if the evidence of specified evidence of the other agent for p is replaced by evidence of unspecified evidence. Perhaps the **EEE Slogan** holds, if it is restricted to evidence of unspecified evidence for p . We would want to include this unspecified evidence, since we assumed (for the purpose of the paper) that the other agent's evidence is true and that this agent uses the same reasoning commitments as we do. How should we revise our credence in p ? The answer depends on the third factor: how well do the pieces of evidence cohere with each other in the light of the proposition p , and how well do they cohere with each other in the light of $\neg p$.

If we had reasons to assume that the third factor is neutral or even positive for the unspecified complementary evidence, then we would be justified in assuming that overall this complementary evidence would increase our credence in p . If we had reasons to assume that it is negative, however, we would not be justified in assuming that the unspecified complementary evidence would increase our credence. The important insight is this: even (higher-order) evidence of unknown, unspecified, complementary evidence is not always evidence. Even assuming **Our Characterization of EE for p**, the **EEE Slogan** is not correct in general.

Indeed, with this observation as a guide, it is easy to construct examples that demonstrate how higher-order evidence of unknown unspecified complementary evidence for some proposition p can be but also can fail to be evidence for p . Let us start with two positive examples.

A More Detailed Milk Example Assume the total evidence of Hanna, tev_{Hanna} , is irrelevant for the question of whether p_1 : there is milk in the fridge. Then Hanna sees Carina searching the fridge for milk and returning empty handed. Thereby Hanna obtains (higher-order) evidence of unknown, unspecified, complementary evidence for $\neg p_1$. That is, that there is some proposition e^* that is included in the total evidence of Carina (i.e. $tev_{Carina} \models e^*$) but not in the total evidence of Hanna (i.e. $tev_{Hanna} \not\models e^*$) such that $\Pr_{Cr_{Carina}}(\neg p_1) = \Pr_R(\neg p_1|e^*) > \Pr_R(\neg p_1|tev_{Hanna}) = \Pr_{Cr_{Hanna}}(\neg p_1)$. Here it is safe to assume that whatever the content of e^* is, the probability of tev_{Hanna} is independent of e^* on the condition of p_1 and $\neg p_1$, and thus that the relevant factor is neutral. In this case the (higher-order) evidence of unspecified, complementary evidence is evidence for $\neg p_1$.

Even stronger support for some proposition p can be obtained by unknown, unspecified comple-

mentary evidence e^* for the proposition in question, if one has reason to assume that it coheres with the rest of one's evidence in the light of p (at least more so than in the light of $\neg p$). For such a case, let us look at a second example.

Lazy Referees Example Suppose two lazy referees need to review the intermediary report of a research project and they decide to divide the whole report in two parts and each assess only one part. The first referee, Silvia, studies the quantity and quality of the published articles, the second referee, Wolfgang, assesses all other criteria for a successful research project. Then they meet again and find out that both their total evidence supports p_1 , the proposition that the research project has been run successfully in the past. Silvia received evidence e and e supports p_1 . Now Silvia learns that Wolfgang obtained some unspecified complementary evidence e^* that also supports p_1 . (Let us assume that all their other evidence is irrelevant to the question of whether or not p_1 .) The fact that they received evidence from diverse areas and that they both came to the conclusion that p_1 leads referee Silvia to expect that her specified evidence e and the unspecified evidence e^* of Wolfgang are coherent in light of the assumption that p_1 , and very incoherent under the assumption that $\neg p_1$.

Thus, in this example it seems appropriate that Silvia considerably increases her credence in p_1 . She has received evidence e for p_1 and higher-order evidence of unknown unspecified evidence e^* for p_1 , and she can assume that only the hypothesis p_1 renders these distinct pieces of evidence (the specified and the unspecified evidence) coherent. Thus, all three factors mentioned in the above **Theorem** are positive, and she should considerably increase her credence in p_1 .

Finally, let us consider a negative case that shows that the **EEE Slogan** is incorrect for some higher-order evidence of unknown, unspecified, complementary evidence; let us consider a variation of Feldman's **Criminal Case Example**:

Another Criminal Case Example Suppose Melina is a detective in a case similar to Feldman's criminal case. Melina believes that she has considered all evidence e that is relevant to the question at hand, and that e points clearly in the direction of p_1 , the proposition that Lefty is the culprit. Melina has refused, however, to consider evidence offered to her by the chief of police, because Melina knows that the chief is engaged in a secret love affair with Lefty. Then suppose the second detective Vincent states that he has collaborated with the chief of police and that they found evidence e^* that incriminates Righty—as opposed to Lefty. Thus, Melina receives higher-order evidence of unspecified, complementary evidence in favour of $\neg p_1$. For Melina it is reasonable to assume that all of this unspecified evidence e^* is true but misleading (the witnesses might be lying, the coroner might be paid off, the police might have planted evidence). Part of Melina's evidence indicates that her evidence e and the unspecified evidence e^* are incoherent given the assumption $\neg p_1$, but coherent given the assumption p_1 . After all, p_1 , that Lefty is the culprit, and the knowledge about the secret love affair predict not only her evidence but also why Vincent found all the (misleading) evidence pointing at Righty, but none of the evidence for p_1 . So in this case higher-order evidence

of unspecified evidence for $\neg p_1$ is not evidence for $\neg p_1$. Melina has reasons to believe that this unspecified evidence does not fit with her evidence, at least not under the assumption that $\neg p_1$. Only p_1 makes the evidence fit together. Since Melina's evidence already supports p_1 anyway, Melina does not revise her credence in p_1 .

This final example shows that evidence of evidence for some proposition is not always evidence for the proposition: the **EEE Slogan** is incorrect for some (higher-order) evidence of unspecified, complementary evidence. These results of course depend on **Our Characterization of EE for p**.

5 Conclusion

The **EEE Slogan** captures one of the most important—albeit vague—ideas in social epistemology. It is an idea we refer to when we defer to the judgements of others who we believe to be better informed than we are. Furthermore, to a large extent the division of labour in scientific inquiry rests on the idea that if we receive evidence that someone else has received evidence for p , this is also evidence for us that p is true. Yet despite its intuitive appeal, and its importance in everyday life and in science, most results so far indicated that the **EEE Slogan** is incorrect.

In a first step, we provided a new framework for understanding and modelling evidence of evidence, which we called **Dyadic Bayesianism**. Within this framework we proposed **Our Characterization of EE for p**. We understand evidence of evidence for p as higher-order evidence for the existence of some (specified or unspecified) unknown evidence in the evidential state of another agent that supports p . In a second step, we investigated under which conditions evidence of evidence for p is evidence for p . Given two idealizing assumptions (i.e. that all evidence is true and that agents share the same reasoning commitments) we argued (i) that evidence of more comprehensive evidence for p is always evidence for p , and (ii) that evidence of complementary evidence for p is not always evidence for p . These results demonstrate the strength of our approach and our framework of **Dyadic Bayesianism** as regards the project to understand evidence of evidence and investigate the **EEE Slogan**.

Acknowledgements Especial thanks go to Mattias Skipper and Asbjørn Steglich-Petersen for detailed commentaries on previous versions of the paper, which helped us to improve it considerably. We are also grateful to Branden Fitelson for fruitful discussions on higher-order evidence that influenced this paper. Anna-Maria A. Eder's research was funded by the Austrian Science Fund (Erwin Schrödinger Program) through the research project *Higher-Order Evidence* (reference number J 3783-G24) at Northeastern University, Boston. Peter Brössel's research on the paper has been generously supported by the German Research Foundation (Emmy-Noether Program) through the research project *From Perception to Belief and Back Again* (reference number BR 5210/1-1.) at the Ruhr-University Bochum.

References

- [1] Adler, J. 2017. Epistemological problems of testimony, In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy (Winter 2017 Edition)*. URL = <<https://plato.stanford.edu/archives/win2017/entries/testimony-episprob/>>.
- [2] Barnett, D. forthcoming. Perceptual justification and the Cartesian theater. In *Oxford Studies in Epistemology* 6.
- [3] Brössel, P. accepted. *Rethinking Bayesian Confirmation Theory*. Springer.
- [4] Brössel, P. 2017. Rational relations between perception and belief: the case of color. *Review of Philosophy and Psychology* 8 (4): 721-741.
- [5] Brössel, P. 2015. Keynes's coefficient of dependence revisited. *Erkenntnis* 80 (3): 521-553.
- [6] Brössel, P. 2013. The problem of measure sensitivity redux. *Philosophy of Science* 80 (3): 378-397.
- [7] Brössel, P. and Eder, AM. A. 2014. How to resolve doxastic disagreement. *Synthese* 19, pp. 2359–81.
- [8] Christensen, D. 2010. Higher-order evidence. *Philosophy and Phenomenological Research* 81. pp. 185–215.
- [9] Comesaña, J. and Tal, E. 2015. Evidence of evidence is evidence (trivially). *Analysis* 75, pp. 557–59.
- [10] Dorst, K. and Fitelson, B. ms. Evidence of evidence: a higher-order approach.
- [11] Eder, AM. A. ms. Evidential probabilities and credences.
- [12] Feldman, R. 2007. Reasonable religious disagreements. In L. Antony (ed.), *Philosophers Without God: Meditations on Atheism and the Secular Life*. Oxford University Press: 194-214.
- [13] Fitelson, B. 2012. Evidence of evidence is not (necessarily) evidence. *Analysis* 72, pp. 85–88.
- [14] Fitelson, B. 2001. *Studies in Bayesian Confirmation Theory*. Ph.D. thesis, University of Wisconsin, Madison.
- [15] Fitelson, B. 1999. The plurality of Bayesian measures of confirmation and the problem of measure sensitivity. *Philosophy of Science* 66, pp. S362–S378.
- [16] Huttegger, S. 2013. In defense of reflection. *Philosophy of Science* 80, pp. 413–33.
- [17] Keynes, J. M. 1921. *A Treatise on Probability*. Macmillan.
- [18] Lackey, J., and Sosa, E. 2006. *The Epistemology of Testimony*. Oxford University Press.
- [19] Lyons, J. 2017. Epistemological problems of perception. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy (Spring 2017 Edition)*, URL = <https://plato.stanford.edu/archives/spr2017/entries/perception-episprob/>
- [20] Moretti, L. 2016. Tal and Comesaña on evidence of evidence. *The Reasoner* 10, pp. 38–39.

- [21] Moretti, L. 2015. Evidence of expert's evidence is evidence. *Episteme* 13, pp. 208–18.
- [22] Pryor, J. 2000. The skeptic and the dogmatist. *Noûs* 34, pp. 517–49.
- [23] Roche, W. 2014. Evidence of evidence is evidence under screening-off. *Episteme* 11, pp. 119–24.
- [24] Shogenji, T. 2003. A condition for transitivity in probabilistic support. *The British Journal for the Philosophy of Science* 54, pp. 613–16.
- [25] Shogenji, T. 1999. Is coherence truth conducive? *Analysis* 59, pp. 338–45.
- [26] Tal, E. and Comesaña, J. 2017. Is evidence of evidence evidence? *Noûs* 51, pp. 95–112.
- [27] Titelbaum, M. forthcoming. *Fundamentals of Bayesian Epistemology*. Oxford University Press.
- [28] Van Fraassen, B. 1984. Belief and the will. *The Journal of Philosophy* 81, pp. 235–56.