

Chapter 5

Judgments About the Relevance of Evidence
in the Context of Peer Disagreements
and Practical Rationality

Amir Konigsberg

Abstract This chapter addresses a core topic in the recent debates about disagreement between peers, namely whether and how you ought to revise your beliefs if you discover that you are disagreeing with a peer – a colleague, fellow expert, or simply someone that you have reason to believe is just as competent as you are on the matter at hand. The topic of disagreement and more specifically the problem of how to respond to disagreement, is relevant in many areas in life where the same information is available to different people that come to hold different beliefs in regard to what that information *means*. The topic also bears on questions relating to epistemic warrant, namely to what extent one’s beliefs can be justified by evidence, to first-person conviction, epistemic humility, normative epistemology, and self-servicing beliefs.

In this chapter I approach this topic differently to how it has been approached in the recent literature. I believe that there is a difference between addressing the disagreement problem hypothetically, or in theory, as an abstraction of *a case of disagreement*, and regarding this same problem when it is considered from a *practical* point of view, in consideration of features that characterize *actual* cases of disagreement. I believe that the core difference between these two cases relates to judgments relating to the relevance of second-order evidence (roughly, evidence about the viability of inferences made from what we regularly treat as evidence). In actual cases of disagreement, as opposed to hypothetical abstractions of such cases, what is or isn’t relevant depends on how uncertain the situation is seen to be by the person in question. And this depends on the level of confidence as perceived by that person’s subjective first-person judgment. In what follows I show that subjective confidence about first-order judgments can swamp second-order evidence – against a plausible view that it shouldn’t. And I believe that this is a significant problem in practical rationality that is brought into focus by disagreement problems. I focus on the problem of *relevancy judgments* and how they relate to the calibration of

A. Konigsberg (✉)
The Center for the Study of Rationality and Interactive Decision Theory, The Hebrew University of Jerusalem, Jerusalem, Israel
e-mail: amirkonigsberg@gmail.com

33 first-person judgments with statistical data, and I demonstrate how this impinges on
34 disagreements when these are considered from the first-person perspective.

35 5.1 Introduction

36 Recent debates dealing with the epistemic significance of peer disagreement have
37 sought to provide responses to cases in which peers disagree about the epistemic
38 import of a shared body of evidence.¹ Various responses have been suggested in the
39 literature.² The problem, as typically addressed in the literature, takes the following
40 general form:

41 If persons A and B are epistemic peers – meaning roughly that it is equally
42 probable that A and B will be correct on matters relating to the domain in which
43 they are peers – and on a particular unexpected occasion they happen to find out that
44 they disagree about whether a particular proposition P (pertaining to that domain) is
45 true given the evidence that is equally available to them both, and assuming that
46 neither party has any independent reason to discount the dissenting party's conclusion,
47 they ought to respond to this discovery . . . *in such and such a way*.

48 The responses to this problem in the literature vary, and can roughly be divided
49 into three types of response: (1) the *bootstrap* response; (2) the *conciliatory*
50 response, and (3) the *egalitarian* response. I think that some of these responses
51 have considerable appeal. But I also think that some make sense theoretically, but
52 do not make sense from a practical point of view, and that others are simply
53 unreasonable. All the same, I do not believe that the goal of establishing *which* of
54 these responses is better, as a good deal of the recent debate has been focused on
55 doing, actually address the important but understated problem of practical rational-
56 ity that underlies the disagreement problem stated above.

57 Customarily, the problem of disagreement asks about the appropriate response
58 (typically the response of a peer) upon encountering a disagreement between peers.
59 The literature offers different solutions to the problem, each of which has more or
60 less normative appeal. Yet none of these solutions seems to engage with what seems
61 to be the *real* problem of disagreement. It is my aim in this chapter to highlight what
62 the real problem of disagreement is. It is, roughly, the problem of deciding *whether*
63 a revisionary tactic is appropriate following the discovery of disagreement, as well
64 as deciding *which* revisionary tactic is appropriate. This non-standard approach to
65 the disagreement problem exposes a slippery and inevitable difficulty that any
66 discussion of disagreement ought to deal with. Once recognized, the real problem

¹ Another subject matter in the literature on disagreement is whether responses to peer disagreement ought to be the same (i.e., should entail the same revisionary response) in all domains. This question is discussed in Konigsberg 2013a.

² Such as the Equal Weight View (Elga 2007), the Total Evidence View (Kelly 2009), the Common Sense View (Enoch 2010) as well as a number of other closely related approaches (Feldman 2006, 2007; Christensen 2007; Matheson 2009; Moss 2011).

of disagreement reflects on the standard question addressed in the literature about
which revisionary tactic is appropriate for dealing rationally with disagreement.
The problem also generalizes to broader problems in practical rationality.

The above-mentioned approaches (1–3 above) are characterized by the different
tactics that they propose for dealing with disagreement. But these tactics only
appear to be relevant after the truly hard work of deciding *whether* they are relevant
in each actual case of disagreement has been done. And this, I believe, is a huge
problem that has not been adequately recognized in the literature or has even largely
been missed until now. It is, in a word, *the problem of judgments about relevancy*.

The epistemic significance of this problem extends beyond debates surrounding
disagreement, since relevancy judgments involve the subjective appreciation of
evidence made about the reliability of inferences made from evidence. It is my aim
here to draw attention to this general problem which, I believe, also lies at the heart
of debates surrounding disagreement. It is my contention that *actual* cases of
disagreement, as opposed to *possible* cases of disagreement, must deal with this
inevitable situation.

The chapter will proceed in two stages. In the next section I outline what I take to
be the real problem of disagreement, setting forth my core argument. But I will start
with some preliminaries. After that I will present three approaches that I take to
characterize the solutions that have been proposed in the literature. In the course of
doing so I will show why most of these do not address the real problem of
disagreement. But I will also suggest which of the approaches in the literature is
most plausible in view of its partial recognition of the underlying difficulty brought
about by the inescapability of relevancy judgments. My foremost aim in this
chapter is to highlight a fundamental difficulty relating to first-person judgments
about the relevance of second-order information – roughly, information which is
applicable to specific cases by virtue of their location in a broader statistical
framework. I aim to demonstrate how this difficulty impinges on actual situations
where revisionary responses to disagreement are called for.

5.2 The Real Problem of Disagreement – *Setting the Stage*

The crux of the matter lies in a practical paradox of sorts, which is inevitable. This
paradox relates to judgments about evidence made from subjective standpoints.
Before I present this practical paradox I will start by clarifying what I mean by
disagreement.

5.2.1 Disagreements Between Ordinary People

Disagreements as I shall refer to them here are situations in which *people* disagree.
More specifically, these types of disagreements involve cases where one person

104 finds out that another person, typically someone whose relevant epistemic capabil-
 105 ities are similar, holds a different opinion, view, or belief on the same matter.³ The
 106 fact that I will be referring to ordinary people in this context is important, and I
 107 make note of it here because it imparts that I wish to relate to the normative question
 108 of how *human* reasoners ought to respond to disagreements in situations that are
 109 characterized by uncertainty about who or what is correct. This approach contrasts
 110 to another, prevalent in the economic as well as the philosophical literature, in
 111 which rational agents rather than human reasoners are the focus.⁴ The rational
 112 behavior of rational agents is typically different to that of human reasoners, and so
 113 as to set the stage for addressing the problem, the distinction must be made.

114 5.2.2 *A Brief Note on Disagreements in the Economic* 115 *Literature*

116 In the economic literature there has been an ongoing debate since the mid-1970s
 117 surrounding the question of disagreement.⁵ The core of this debate focuses on the
 118 possibility of rational disagreement between rational agents. The question posed is
 119 whether it is possible that agents who are expected to conditionalize on information
 120 in the same way, can agree to disagree (Aumann 1976). The debate considers
 121 whether it is possible for rational agents to disagree *rationally*.

122 In the cases of disagreement referred to in the economic literature the types of
 123 agents referred to are *not* human agents. And the rationality that is attributed to
 124 these agents is perhaps *not* human rationality. Human rationality, in the context of
 125 disagreement, relates to human reasoners that encounter evidence to which they
 126 know that they may respond imperfectly. The problem of disagreement is thus
 127 located in the wider context of human fallibility and regards disagreements as
 128 opportunities for corrective measures aimed at mitigating erroneous consequences
 129 of imperfect reasoning.

³ While opinions, views, and beliefs may suggest different meanings, each suited more than the other for a particular context; I use them here interchangeably as referring to what a person regards as true.

⁴ By rational agents I have in mind something similar to what Thaler and Sunstein have recently referred to as *Econs*. See (Thaler and Sunstein 2008).

⁵ See for instance (Aumann 1976) and (Geanakoplos and Polemarchakis 1982; Cave 1983; Moses and Nachum 1990; Rubinstein and Wolinsky 1990; Hanson 2003; Dégrement and Roy 2009; Hansen and Cowen 2004; Milgrom and Stokey 1982).

5.2.3 Human Imperfection and Its Implications for Practical Reasoning Problems 130 131

There are two critical senses in which human imperfection impinges on disagreements. The first relates to the gathering of information, the second to inferring conclusions from it. People’s capacities in both these areas are limited, and they commonly make mistakes. In the course of my discussion I will assume that people are typically aware that they make mistakes – that we know that we are not normally capable of taking all or perhaps unlimited information into account when we deliberate in uncertain conditions. Moreover, we are normally also aware of the fact that when we do possess information that is relevant to our beliefs, our responses – characteristically the inferences that we make from this information – are often imperfect. And by ‘imperfect’ I have in mind, approximately, three things. 132 133 134 135 136 137 138 139 140 141

5.2.4 Imperfection and Reasoning 142

Firstly, when people infer conclusions from evidence, however limited or encompassing this evidence may be, they are not always correct in what they infer. And people commonly recognize this about themselves. What this actually means is that they recognize that their reasoning is error prone, and thus imperfect. Secondly, while people generally know that their reasoning is error prone, they do not always recognize the occasions in which it is so. Because of this, people often think they are right when they are wrong, and therefore incorrect reasoning sometimes goes unrecognized. Thirdly, because they know that they may sometimes be wrong about what they believe, and because they also know that they do not always recognize the occasions when this is so, people should not always be certain that what they believe is correct. Indeed, it appears that in general we ought to have some reservations about the viability of our responses, particularly when they encounter dissent from an esteemed counterpart or a fellow expert. 143 144 145 146 147 148 149 150 151 152 153 154 155 156

5.2.5 When Subjective Credence Plays Against Facts of the Matter 157 158

In many situations of uncertainty, evidence may be more or less convincing, and this seems to play subjective credences against perhaps unknown facts of the matter. All the human reasoner has to go by is his subjective credence, which is assumed to mirror the strength of the evidence that he has. In cases such as this, an individual’s subjective flaws as an evidence evaluator prescribe some type of risk mitigating strategy so that inferences whose impact on credence is partly subjective, can be weighed against some type of objective standard that is not based on the same error prone reasoning. 159 160 161 162 163 164 165 166

167 **5.2.6 First-Order and Second-Order Evidence**

168 A helpful distinction has been made in the literature between two kinds of evi-
 169 dence.⁶ The distinction provides a convenient taxonomy for considering corrective
 170 standards for mitigating erroneous tendencies in reasoning. The distinction is
 171 between first-order evidence and second-order evidence and it can be understood
 172 as making a point about two kinds of epistemically relevant considerations, or
 173 evidence.

174 First-order evidence refers to the kind of evidence the presence of which can
 175 increase or decrease subjective credence in a particular proposition. And by
 176 increase, I mean that it can make one more confident about the proposition than
 177 before – either by moving one's belief towards that expressed by the proposition, or
 178 by strengthening one's belief. In the same way, first-order evidence can decrease
 179 one's confidence too. Moreover, if first-order evidence is epistemically insignifi-
 180 cant, it may neither increase nor decrease credence.

181 The notion of second-order evidence refers to evidence that bears on whether
 182 one's beliefs, or the credences that one has assigned, are likely to be correct.⁷ For
 183 our purposes what is entailed by first-order evidence is partly subjective because
 184 credences that are based on first-order evidence are estimates of a proposition's
 185 truth value, based on subjective assessments of first-order evidence. Credences
 186 might appear precise because they are usually expressed in numerical form, but
 187 let's remember that credences are people's estimates of truth values, expressed as
 188 probabilities. As opposed to first-order evidence, which usually relates to a person's
 189 present judgment, second-order evidence is typically evidence that is based on past
 190 epistemic performance, or experience, and as such does not depend on corrobora-
 191 tion by present judgment.

192 Here is an example illustrating this distinction. The first-order evidence (FOE)
 193 that I encounter may be the Candlestick in the Hall, which supports my belief that
 194 Colonel Mustard did it (P): $\langle \text{FOE} \mid P \rangle$, or else it increases credence in the belief
 195 that he did it. The second-order evidence (SOE) that I possess may be prior
 196 knowledge that in the past, when I inferred *who* was guilty on the basis of weapon
 197 and location alone (FOE), I was wrong 70 % of the time. In this case second-order
 198 evidence is the knowledge that I have about my past performance in inferring
 199 conclusions in similar conditions (using the same variables).⁸ It tells me how likely
 200 it is that my inference – e.g., $\langle \text{FOE} \mid P \rangle$, is correct.

201 In the context of disagreements between peers, first-order evidence will be the
 202 evidence that each peer encounters and which consequently leads each to believe as

⁶ Kelly (2009), Christensen (2007), and Feldman (2006) also refer to this distinction.

⁷ Compare Kelly's discussion of Downward Epistemic Push (Kelly 2009, sec. 5.3).

⁸ Second-order evidence, conceived as I am presently suggesting, provides information about the likelihood of some event, outcome or possibility in some general population of events, outcomes, or possibilities. It may take many different forms and prior experience or performance is only one such form.

he does in the first-place. Therefore if the disagreement is, for instance, between weather forecasters, and concerns the weather forecast for tomorrow, first-order evidence is the evidence on which each person bases his belief about tomorrow's weather. More generally, we might think of this as the type of evidence that is normally needed for a weather forecaster to make up his mind about tomorrow's weather, prior, that is, to finding out what his peer believes about it. Thus first-order evidence in this context may perhaps be temperature maps (TM), atmospheric factors (AF), and other metrological features (MF) on which weather forecasters typically base their predictions. Let E denote a particular piece of evidence. In this case a prediction based on FOE will look like this: $\langle \text{ETM}; \text{EAF}; \text{EMF} \mid \text{P} \rangle$.

Second-order evidence will typically be evidence that relates to the belief-forming circumstances in which conclusions are inferred from first-order evidence. This may for instance include considerations about how likely it is that the inferences made from first-order evidence are correct. In disagreement problems, second-order evidence will characteristically relate to prior knowledge about a person's competence in inferring conclusions from first-order evidence (typically stated in probabilities based on prior performance), or perhaps knowledge about the person's susceptibility to error.

The widely consensual position in the literature is that in reasoning problems that involve first and second-order evidence, subjective credences that are based on first-order evidence ought to be balanced by probabilities derived from second-order evidence based on past performance (in the same way circumstantial indicators in Bayesian reasoning problems are weighed against base-rate information). The general contention is that not to take into account second-order evidence, usually prior probabilities, where these are informative and thus epistemically relevant to the assessment of the viability of present evidence, is a failure of reasoning. To be more specific, it is a failure to consider objective – i.e., second-order as well as subjective – i.e., first-order – factors, both of which are of epistemic value. Otherwise put, to only consider first-order evidence and not to consider second-order evidence is a failure to consider *all* the relevant evidence.

This insight has been expressed in the literature. Kelly, for instance, articulates it as follows: “what it is reasonable to believe about the world on the basis of one's evidence is constrained by what it is reasonable to believe *about* one's evidence” (Kelly 2009, sec. 5.3). So too, Christensen notes that “the rationality of first-order beliefs cannot in general be divorced from the rationality of certain second-order beliefs that bear on the epistemic status of those first-order beliefs” (Christensen 2007, 18). To apply this insight to our previous example, this would mean that if, after encountering the Candlestick in the Hall I infer that Colonel Mustard did it, without considering that second-order evidence suggests that my inference that Colonel Mustard did it is 70 % likely to be wrong, I would be neglecting relevant and thus epistemically valuable evidence.

244 **5.2.7 Theoretical and Practical Prescriptions**

245 The prescription to consider second-order evidence in conjunction with first-order
 246 evidence appears to be unproblematic in theory. In situations of uncertainty,
 247 second-order evidence – usually prior probabilities – should be weighed against
 248 subjective likelihoods derived from first-order evidence, typically in accordance
 249 with Bayes' rule.⁹ But while in theory this appears unproblematic and straightfor-
 250 ward, there seems to be a difficulty in realizing this prescription in practice.

251 Because the normative prescription is that second-order evidence ought to be
 252 incorporated in judgment in situations of uncertainty, it is subjective judgment that
 253 is responsible for determining a situation *as* uncertain, and consequently for
 254 determining *whether* second-order evidence is relevant as a risk-mitigating measure
 255 for addressing this uncertainty. But because determining whether a situation is
 256 uncertain depends, at least in part, on how confident the person is about first-order
 257 evidence, it can make the subjective judgment about *whether* second-order evi-
 258 dence is relevant dependent on the very unlawful outcome it is there to mitigate.
 259 And this is what paves the way to the real problem that a person that encounters
 260 first-order evidence appears to face.

261 The reason this is a practical rather than a theoretical problem is because in
 262 theory there is no apparent difficulty of incorporating second-order evidence in
 263 judgments about first-order evidence; in these kinds of theoretical cases, any
 264 judgment based on first-order evidence will be weighed against the second-order
 265 evidence that applies to that judgment. Why? Well because this is the normative
 266 thing to do; second-order evidence *is* the means by which an individual judgment is
 267 put in the context of passed judgments and thus undergoes statistical corroboration.
 268 It is what is needed so that the viability of an individual judgment can be properly
 269 assessed. But in practice, whether or not second-order evidence applies to an
 270 individual judgment depends on how uncertain the person is or how compelling
 271 first-order evidence appears to be. If first-order evidence is compelling, second-
 272 order evidence may be taken to be irrelevant or inapplicable to that judgment
 273 because statistical corroboration appears unnecessary or inappropriate. Hence
 274 second-order evidence can be regarded as more or less applicable depending on
 275 the diagnostic value of individual judgment, which it was initially aimed to
 276 corroborate.

277 The normative problem that I am trying to outline is a problem of disagreement
 278 between peers considered from a practical point of view. And it is importantly
 279 different to theoretical abstractions of such disagreements. In theory, if you and I
 280 are peers and we disagree (even though each of us is confident about what we

⁹ Bayes's rule, or theorem, is a rule for operating on numerically expressed probabilities to revise a prior probability (in other words, the base-rate) into a posterior probability after new data have been observed. According to the theorem, the posterior probability for event H1 *after* data D is observed and accounted for is: $p(H1|D) = p(H1) p(D|H1)/p(D)$, where $p(H1)$ is the prior probability assigned to H1 *before* D is observed.

believe), second-order evidence about our peerhood is the means by which we locate ourselves in a statistical framework. The fact that we may find first-order evidence compelling seems to be neither here nor there; because disagreement exists between the two of us, and we are equally competent reasoners, second-order evidence must be used, since second-order evidence will have to be judged relevant to the case at hand in some sense and to some degree. In this case the question that remains is how to deal with first-order and second-order evidence in relation to each other. And this is what has traditionally been debated in the literature. But in practice, if you regard first-order evidence as sufficiently compelling, you will probably not regard the evidential situation as uncertain. And as such, the corroboration of your belief with second-order evidence may seem irrelevant or even damaging to your subjective evaluation. In such a case, the fact that your peer disagrees may be evidence in favor of their being wrong. But it is important to see that even if, in such a case, you do not regard your situation as uncertain this does not mean that you are denying that there is a case of peer disagreement. You're not. Peer disagreement can still be a problem even if you yourself are confident about the first-order evidence. I mention this here because one might object that if you find that there is no uncertainty, then this is in fact a denial that there is a case of peer disagreement. But this only seems to be the case if one assumes that an agent's uncertainty about first-order evidence alone merits his using second-order evidence to corroborate his belief.

5.2.8 Practical Problems with Theoretical Prescriptions in Disagreement Problems

The practical difficulty with implementing the normative prescription to weigh first-order evidence against second-order evidence in situations of uncertainty is that from the first-person standpoint second-order evidence often has ambiguous implications. To see this, consider a hypothetical situation. Assume that I know that based on past performance I am 70 % likely to be correct in my predictions. On first thought, I can take this to mean that there is a 70 % chance that my next prediction will be correct and a 30 % chance that it won't. Now, having made my next prediction there appears to be no way for me to ascertain, independent of relying on my present judgment and the various considerations that support it, whether my prediction falls in the positive or negative percentiles of chance. That is to say, I have no way of knowing whether my present belief is an instance affirming the 70 % chance that I am correct or the 30 % that I am not. My probability of being correct, based on past performance, is second-order evidence the inclusion of which appears to depend on my present level of confidence about first-order evidence. And if I my confidence is high, I may regard my present judgment as being an affirming instance of the positive likelihood of my being correct (in line with how second-order evidence can be understood), rather than an uncertainty in need of

321 corroboration. And because, on this interpretation, *I am* correct, there is no need to
322 weigh my present level of confidence against second-order evidence.¹⁰ Note that I
323 am not suggesting that I would necessarily be justified in neglecting the 70:30 base-
324 rate. Rather, what I am suggesting is that the significance of the 70:30 base rate can
325 mean different things in terms of epistemic justification.

326 This suggests that aside from the normative prescription to weigh first-order
327 evidence against second-order evidence I appear to also have a normative obliga-
328 tion to treat first-order evidence in accordance with the level of epistemic warrant
329 that it provides. It does after all seem to be uncontroversial that different evidential
330 situations warrant varying levels of confidence. And it seems that proper incorpo-
331 ration of new information about these situations depends on the epistemic warrant
332 that is provided by first-order evidence. A person in a first-person standpoint may
333 consequently be faced with two, possibly conflicting, normative prescriptions, from
334 which the suspected practical paradox arises:

- 335 (1) Respond to first-order evidence in accordance with how convincing it appears
336 to be.
- 337 (2) Mitigate the risk of being wrong by weighing first-order evidence against
338 second-order evidence.

339 As noted, the practical problem here is that there is no independent way to
340 ascertain which prescription – (1) or (2) – applies, and there is no immediately
341 obvious or straightforward weigh of combining them. And this situation, in which
342 we are asked about an individual's appropriate response once conflict with a peer is
343 discovered, appears to be reflective of a class of epistemically ambiguous situations
344 in which the crux seems to lie in an individual's ability to determine the appropriate
345 revisionary response to the situation. And because doing so is largely a matter of
346 how confident the individual is about the first-order evidence, the inclusion of
347 second-order evidence seems to depend, at least in part, on the selfsame risk-
348 prone reasoning it is there to mitigate. And this, it seems, is inevitable. As long as
349 subjective judgment is responsible for deciding *whether* second-order evidence is
350 relevant, it doesn't appear to matter that second-order evidence is independent of
351 current judgment, the risk of fallible subjective judgments remains.

352 5.2.9 *The Real Problem of Disagreement – A Practical* 353 *Paradox*

354 If the normative prescription is that an individual in a decision circumstance ought
355 to decide whether or not second-order evidence is relevant to that circumstance, he

¹⁰ Kelly states of cases such as these that “one’s first-order evidence not only confirms the belief in question; it also confirms a proposition to the effect that it is reasonable for one to hold that belief” (Kelly 2009, sec. 5.3).

must have the ability to distinguish when it is and when it isn't relevant. We assume that this is a function of how ambiguous the evidence is, perhaps how weak the evidence is (Enoch 2010; Kelly 2005), or how uncertain he is about what he has inferred from the evidence. In each case it is on the basis of such considerations that the individual deems second-order evidence relevant or irrelevant. And this seems to lead to a situation where the judgment of relevance has no independent evidence to go by. Hence theoretically, Bayes rule may well offer a precise way to weigh beliefs, and philosophers writing about disagreement may well make suggestions about which responses are appropriate for peers that encounter conflicting beliefs. But tactics such as these only seem to be relevant *after* the hard work of deciding relevancy has been done, and this, as I have tried to show, is theoretically an underdescribed problem that I believe has largely been missed until now.¹¹

5.3 Responses in the Literature

I believe that responses to peer disagreement in the recent philosophical literature can be divided into three kinds. In the next section I will address each of the approaches in the literature in relation to the problem highlighted in Sect. 5.2. Before doing so I will sketch a pseudo-particularized example of disagreement on the basis of which the plausibility of each of the approaches can be assessed.

5.3.1 A Case of Peer Disagreement

Jill and Jack are two equally ranked chess masters. As it happens, Jill and Jack have other things in common aside from sharing the same title and rank at chess. They have both been playing chess for the same number of years and they have won the same number of games, at equally ranked tournaments, against equally classed players, using similar game strategies. Additionally, Jill and Jack also know all of this about each other.

On a particular occasion, Jill and Jack are each independently asked by an examiner which color has the advantage in a particular chessboard arrangement. Jill tells the examiner that she thinks that White has the advantage; Jack tells the examiner that he thinks that Black has the advantage. Then each of them is told by the examiner about what the other thinks. What should Jill and Jack do in regard to their beliefs after being given this information, assuming that is, that neither one of

¹¹ Elsewhere I discuss the epistemic significance of relevancy judgments. In this context see also Maya Bar-Hillel's seminal "The base rate fallacy in probability judgments" (Bar-Hillel 1980), which focuses on relevancy judgments in establishing whether or not base rates ought to be incorporated in probability judgments. See also: (Bar-Hillel and Fischhoff 1981; Bar-Hillel 1982; Welsh and Navarro 2012; Barbey and Sloman 2007; Ajzen 1977).