

A Game-Theoretic Approach To Peer Disagreement*

Remco Heesen[†] Pieter van der Kolk[‡]

December 21, 2015

Abstract

In this paper we propose and analyze a game-theoretic model of the epistemology of peer disagreement. In this model, the peers' rationality is evaluated in terms of their probability of ending the disagreement with a true belief. We find that different strategies—in particular, one based on the Steadfast View and one based on the Conciliatory View—are rational depending on the truth-sensitivity of the individuals involved in the disagreement. Interestingly, the Steadfast and the Conciliatory Views can even be rational simultaneously in some circumstances. We tentatively provide some reasons to favor the

*This paper has been published in *Erkenntnis*. The final publication is available under open access at Springer via <http://dx.doi.org/10.1007/s10670-015-9800-8>. Thanks to Kevin Zollman, Jan-Willem Romeijn, Frank Hindriks, Liam Bright, Nathaneal Smith, Lauren Leydon-Hardy, Matt Frise, two anonymous reviewers, and audiences at the 2014 meeting of the Central States Philosophical Association and the 2014 Rochester Graduate Epistemology Conference, including in particular Laurie Paul, Sandy Goldberg, and Earl Conee, for valuable comments.

[†]Department of Philosophy, Baker Hall 161, Carnegie Mellon University, Pittsburgh, PA 15213-3890, USA. Email: rheesen@cmu.edu

[‡]Faculty of Philosophy, University of Groningen, Oude Boteringestraat 52, 9712 GL Groningen, The Netherlands. Email: p.m.van.der.kolk@rug.nl

Conciliatory View in such cases. We argue that the game-theoretic perspective is a fruitful one in this debate, and this fruitfulness has not been exhausted by the present paper.

1 Introduction

The aim of this paper is to show that the problem of peer disagreement can be analyzed from a game-theoretic perspective. The problem of peer disagreement, as it is presented in the literature (e.g., Kelly 2005, 167; Christensen 2009, 756; Elga 2007, 478; Feldman 2007, 201), is how to respond rationally to the disagreement from an epistemic peer, whereby *epistemic peer* is construed as an agent who has the same evidence and is comparably good at evaluating that evidence (Kelly 2005, 170; Christensen 2007, 188; Feldman 2007, 201; Lackey 2008, 274). Game theory, in turn, is the study of strategic decision making, where ‘strategic’ means that the decision of one decision maker may interact with that of another. This paper explains how the latter can be used to analyze the former.

To do so, we focus on two prominent strategies recommended in the literature about peer disagreement, namely the response advocated by the *Conciliatory View* and the one suggested by the *Steadfast View*.¹ On the Conciliatory View, it can *not* be rational for an agent to stick to her opinion when it is disputed by an epistemic peer. Instead, she should suspend judgment (Feldman 2007), split the difference (Elga 2007), or at least migrate her opinion significantly in the direction of her peer’s conflicting opinion (Christensen 2007). In this paper we focus on full belief states rather than degrees of belief, so that the subtle differences between these ‘Conciliatory Views’ can be dispensed with. According to the Steadfast View, on the other hand, it *can* be rational for an agent to retain her opinion in the face of peer

¹In the debate about peer disagreement, it is common to talk about ‘responses’, whereas in the context of game theory ‘strategies’ is conventional. In this paper, we will use the two terms interchangeably.

disagreement (Kelly 2005, van Inwagen 2010).

The game-theoretic toolkit enables us to analyze the rationality of these responses (strategies) for disagreeing peers (players), relative to these peers' epistemic goals (preferences). In the literature on peer disagreement, the epistemic goal is commonly understood to be believing the correct truth-value of the proposition under discussion (Christensen 2007, 216; Feldman 2007, 212; Elga 2007, 488; Kelly 2010, 17; and, even if only indirectly, White 2005, 450). Thus, the rationality of the available responses—i.e., the Conciliatory strategy and the Steadfast strategy—can be analyzed by investigating to what extent they satisfy the preferences (epistemic goals) of the disagreeing peers. In section 2 we argue that existing formal approaches do not address this particular question. In sections 3 and 4 we explain the details of our approach to the problem of peer disagreement. Section 5 discusses the results of this model, section 6 considers some possible extensions or variations of the model, and section 7 wraps up by emphasizing some key take-aways.

2 Why a Game-Theoretic Approach?

Why should a game-theoretic analysis be a relevant contribution to the debate about peer disagreement? Our motivation is that the resources of game theory enable a clarification of the responses to peer disagreement—in particular, of the Conciliatory View and the Steadfast View—along an independently motivated and well-developed standard. In the debate about peer disagreement, it is not always clear how exactly rationality is understood, what exactly counts as a peer, what a disagreement is, or even what exactly the Conciliatory View and the Steadfast View amount to (cf. Jehle and Fitelson 2009, Moss 2011, Lasonen-Aarnio 2013).

A formalization along the lines of game theory forces us to be precise about these notions. And the fruit of such explicitness is that it helps us to gain a better understanding of the conditions under which a particular strategy (like the ones suggested by the Conciliatory View and the Steadfast

View) can be considered a rational response to the disagreement from a peer.

We do not want to suggest that our game-theoretic model is the only way to make the machinery under the problem of peer disagreement formally precise. Here we consider some previous work along these lines.

First it is important to distinguish quantitative and qualitative cases of peer disagreement. In *the quantitative case* the agents assign different degrees of belief to a proposition, whereas *the qualitative case* concerns full belief states (belief, disbelief, and suspension of judgment). Some have argued that the quantitative model of epistemic agents should be taken as basic and the qualitative model should be reduced to it (Lin and Kelly 2012, Leitgeb 2014). Others have argued the reverse (Easwaran forthcoming). This debate remains unresolved. As a result, we can treat quantitative and qualitative cases of peer disagreement as separate problems. Our focus in this paper is on the qualitative case. But as the majority of the work in formal epistemology that is potentially relevant to peer disagreement focuses on the quantitative case, we discuss this work first.

There are two dominant models in the literature on revising degrees of belief in light of new information (here, the information that an epistemic peer assigns different degrees of belief). One is the (iterated) *linear pooling* model developed by French (1956), DeGroot (1974) and Lehrer and Wagner (1981). In this model, the revised degrees of belief are obtained by taking a weighted average of the agents' opinions. This is consistent with both the Steadfast View and the Conciliatory View. The Steadfast View says an agent can rationally give weight one to her own opinion and zero to her peer's, whereas the Conciliatory View says this is not rationally permissible.²

However, it is not clear what gives linear pooling its normative force. Without an interpretation of the weights used, "it is not clear why we should change our beliefs according to the weighted linear average, instead of, for

²On this explication of the views French and DeGroot are proponents of the Steadfast View, and Lehrer and Wagner are proponents of the Conciliatory View. However, Lehrer and Wagner would not endorse stronger interpretations of the Conciliatory View, e.g., that giving equal weight to each agent's opinion is required (cf. Jehle and Fitelson 2009).

instance, the weighted geometric average” (Martini et al. 2013, 887).

Romeijn (2015) attempts to give such an interpretation. He shows that if the agents’ priors take a particular form, linear pooling can be construed as a special case of *Bayesian conditionalization* (the other dominant model for revising degrees of belief), where the weights assigned to agents are identified with the truth-conduciveness of those agents. On this construal, linear pooling inherits the normative force that Bayesian conditionalization is generally taken to have, although particular assumptions need to be in place in order for linear pooling to be sanctioned by the Bayesian model.

Two problems remain. First, there appears to be no normative reason for the agents’ priors to take the required form. Second, it does not settle the debate between the Steadfast and the Conciliatory View, as the formalism itself does not settle whether an agent is rationally permitted to give weight one to her own opinion.

The first problem can be circumvented by allowing the agents to have any priors, taking Bayesian conditionalization as the normative model for revising degrees of belief without requiring that it agrees with linear pooling. Under certain assumptions, the agents can be guaranteed to reach a consensus in this model (Aumann 1976, Geanakoplos and Polemarchakis 1982). But the second problem remains.³

It appears, then, that none of the extant work in formal epistemology yields a view on the quantitative case of peer disagreement, although a focused discussion of the relations between the models we discussed and peer disagreement may still yield valuable insight. While we offer no view on

³Geanakoplos and Polemarchakis (1982, 197) show that one of the agents’ opinion may never change despite the presence of disagreement. This suggests that sticking to one’s opinion is (sometimes) rationally permissible, in support of the Steadfast View. But this argument works only if one assumes that Bayesian conditionalization is the *only* requirement of diachronic rationality, since further requirements of rationality may rule out these cases. Moreover, the kind of cases considered by Geanakoplos and Polemarchakis arguably do not count as cases of peer disagreement strictly speaking, as the agents have different evidence concerning the proposition they disagree about.

the quantitative case here, the model we present could relatively easily be adapted to it.

In addition to the problems mentioned above, linear pooling and Bayesian conditionalization offer no solution to the qualitative case of peer disagreement, which will be our focus from here on out. For the qualitative case there are again two dominant classes of relevant formal models. The first is known as *belief revision*, usually (but not necessarily) using the so-called AGM model (Alchourrón et al. 1985). This model has been applied to peer disagreement (Cevolani 2014, Elkin 2015). While these papers are interesting, they beg the question in favor of the Conciliatory View: they explore ways in which a Conciliatory response to peer disagreement affects an agent's other beliefs.

A similar problem holds for the second class of models, those based on *judgment aggregation*. Regardless of whether one follows the dominant axiomatic approach (List and Pettit 2002, List 2013) or focuses more directly on the reliability of aggregation methods (Hartmann et al. 2010, Hartmann and Sprenger 2012), these models already assume that one has decided to form a consensus opinion. Again, the Steadfast View is ruled out by the formal setup without argument.

It is also worth noting that most models of judgment aggregation and voting theory more generally concentrate on the case of at least three agents, whereas we, following the peer disagreement literature, focus on the case of two agents. Most of the prominent aggregation methods rely on some variation of majority voting, which does not yield very interesting results in the case of two disagreeing agents. We briefly return to the case of more than two agents in section 6.

The model of the present paper addresses the peer disagreement debate head on, as we give a direct comparison of the Conciliatory and the Steadfast View. While some previous work has aimed to make ideas from the peer disagreement literature formally precise (Jehle and Fitelson 2009, Cevolani 2014, Elkin 2015), we are not aware of any formal work that makes this kind

of direct comparison.⁴ Some of the work mentioned above could perhaps be adapted to make such a comparison, which we think would be very interesting. But in the remainder of this paper we aim to argue (1) that the specific game-theoretic model we provide captures one interesting way to make the ideas underlying the peer disagreement debate more precise, and (2) that the model is flexible enough that it can be straightforwardly adapted to capture other ways of making these ideas more precise.

3 The Peer Disagreement Game

We introduce our game-theoretic setup with the help of an informal example. Imagine two detectives, call them Jane (Marple) and Hercule (Poirot), who both have been asked to go to a crime scene to investigate whether ϕ , say, whether the butler is the culprit. We make the following three assumptions about the detectives. First, they have the same evidence at their disposal to investigate ϕ , namely whatever traces are left at the crime scene. Second, the detectives can make an informed estimation of how reliable each of them is in investigating ϕ , based on their respective *track-records*; the number of crimes they have solved in the past compared to the number of crimes they did not solve. Third, the detectives really want to find out the truth regarding ϕ , they really want to solve the case.

We take it that the fulfillment of these three conditions is what is (at minimum) required for the two detectives to be called each other's peers, considering the construals of peerhood by, for example, Kelly (2005, 175), Elga (2007, 484), Lackey (2008, 274), and Christensen (2009, 757). The attribution of peerhood then depends on how equal the detectives must be in their reliability. Our analysis accommodates this.

⁴Except perhaps Lehrer and Wagner (1981), as they argue for normative claims which entail that agents should give positive weight to their peers. This is effectively an argument for the Conciliatory View in the quantitative case (at least on liberal interpretations of that view, see footnote 2). But note the criticism by Martini et al. (2013) mentioned above.

Jane and Hercule both go to the crime scene, and spend some time examining and evaluating the evidence. After some time, they meet up to report their findings.

Two things can happen at this point. Jane and Hercule have either formed the same belief about ϕ , or they have formed conflicting beliefs and disagree about ϕ . In the model, these beliefs are generated probabilistically (see the next section).

If the detectives have reached the same conclusion about ϕ , say, they agree that the butler is indeed the culprit, then there is no problem of peer disagreement. The detectives can go write their reports. The case that we are interested in is when the detectives have formed conflicting opinions regarding ϕ ; for example, when Jane believes that the butler is the culprit and Hercule believes that the butler is innocent. And our question is what, in such a case, a rational response for Jane and Hercule can be, given their goal of finding out the truth about ϕ , and the information they have about each other's track-records.

Based on the debate about peer disagreement, we distinguish three strategies that the detectives can choose. The first comes from the Steadfast View and is the strategy of staying with the initial belief. We call this strategy **Stay**. The second strategy is the Conciliatory View's recommendation to suspend judgment.⁵ This strategy is called **Suspend**. And third, for the sake of completeness, we include switching to the belief of the other detective as a third possible strategy, called **Switch**.

After Jane and Hercule find out that they disagree about whether the butler is the culprit, they each play one of these three strategies. When Jane plays **Suspend**, she withdraws her initial belief about ϕ , goes back to the crime scene to re-examine the evidence, and forms a new belief about ϕ . But when Jane plays **Stay**, she chooses to ignore the disagreement and maintains

⁵Technically, the recommendation from the Conciliatory View can also be to split the difference with one's peer (Elga 2007), or to revise one's initial confidence level in the proposition considerably (Christensen 2007). But since we restrict ourselves to full beliefs, we take it that the Conciliatory View's recommendation amounts to suspending judgment.

her initial opinion. And when Jane plays **Switch**, she chooses to ignore her own opinion and takes over the belief of Hercule.

So only when a detective plays **Suspend** she gets a chance to form a new opinion. It might be objected that acquiring a new belief is not a necessary consequence of suspending judgment. We agree. We should distinguish between two ways in which judgment can be suspended. The first is to suspend judgment *indefinitely*, or at least until new evidence comes in, because there is at present not enough evidence to form a rational belief. The second is to suspend judgment only *momentarily*, as an act of caution in light of unexpected counterevidence, but after which a new belief may be formed through a re-examination of the evidence. Such a momentary suspension of judgment is justified for cases in which a Peircean ‘irritation of doubt’ needs to be resolved, because it is unsatisfactory or unwarranted not to have a belief about the matter. We take it that this is the preferred form of suspension of judgment in the well-known restaurant case of Christensen (2007, 193), in which two peers disagree about the division of the bill, as well as in other influential examples of peer disagreement (e.g. Feldman 2007, 208–209). In this paper we also work with this short-term interpretation of suspension of judgment. A long-term interpretation would be a welcome extension of our analysis (see section 6 and appendix B).

The disagreement game ends when the two detectives reach an agreement about ϕ . For example, suppose Jane believes that the butler did it, and Hercule believes that he did not do it, and Jane plays **Suspend** and Hercule plays **Stay**. Then the game ends when, after re-examining the evidence, Jane draws the same conclusion as Hercule, namely that the butler is innocent.⁶ The same would happen when, for example, Jane plays **Stay** and Hercule plays **Switch**. But the game continues when, after one or both of them re-examine the evidence, the two detectives still disagree about ϕ .

⁶Since Hercule plays **Stay**, he never changes his belief. So the game ends when Jane concedes. As we explain in more detail in section 4, our probabilistic model for generating new beliefs guarantees that this will happen eventually when she plays **Suspend**.

For the purposes of this paper, we assume that the detectives do not change strategies throughout the disagreement game.⁷ This means that the game might also continue forever. For example, when Jane believes that the butler is innocent and Hercule disagrees, and both detectives play **Stay**, then they will never come to an agreement. The same thing happens when both detectives play **Switch**.⁸

And now we are in a position to analyze how well these strategies do in guiding each detective to the correct verdict on whether the butler did it. Which of these strategies gives a detective the best prospects of arriving at the truth?

Observe that which strategy is best will depend on two factors.⁹ First, it depends on the reliability (i.e., the track-record) of each of the two detectives. For example, if Jane thinks that Hercule is better at evaluating correctly whether the butler did it, then it would be ill-advised for her to play **Stay** upon finding out that Hercule disagrees with her initial assessment. But when Jane thinks that she is more reliable than Hercule, then playing **Stay** may be sensible.

Second, which strategy is best depends also on the strategy of the other detective. For example, when Hercule plays **Stay**, it does not really matter for Jane whether she plays **Suspend** or **Switch**, because either way the game

⁷The reason is that this allows a straightforward comparison of the Conciliatory View, which recommends playing **Suspend** for all instances of peer disagreement, and the Steadfast View, according to which playing **Stay** can be rational. It would be an interesting extension of our model to allow players to change their strategy during the game (see section 6).

⁸That under these strategies the game continues forever does not make an evaluation of the rationality of these strategies impossible. For in both cases we can still evaluate how well these strategies do with respect to tracking the truth.

⁹It should be noted that on our approach the rationality of a strategy does not depend in any way on ‘right reasoning’ at the first stage, during the initial assessment of the evidence, like it does in Kelly (2005, 2010). Our approach is more akin to Christensen (2007) or Elga (2007), where a rational strategy is to be determined independent of one’s initial reasoning behind the disputed belief.

will end when Jane takes over the conclusion of Hercule. But when Hercule plays **Switch**, it *does* matter whether Jane plays **Suspend** or **Switch**, because playing **Switch** will bring them in a state of perpetual disagreement, whereas playing **Suspend** will make them agree eventually (due to the probabilistic way in which new beliefs are generated; see the next section). We will return to these points in section 5.

This concludes our informal description of the peer disagreement game. In the next section we will provide the formal vocabulary, and then analyze this game.

4 Rationality for Jane and Hercule

Whenever Jane and Hercule investigate the evidence, they may conclude that the butler did it (ϕ) or that he did not do it ($\neg\phi$). One of these conclusions is *true* and one is *false*.

We will denote by p and q the reliability or *truth-sensitivity* of Jane and Hercule, respectively. Thus p is the probability, on any given investigation, that Jane draws a true conclusion from the evidence. $1 - p$ denotes the probability of a false belief. So if the butler really did it Jane believes that he did it with probability p and believes in his innocence with probability $1 - p$. Whereas if he is innocent she believes in his innocence with probability p and believes that he did it with probability $1 - p$. Hercule's probabilities of drawing a true or a false conclusion from the evidence are denoted by q and $1 - q$, respectively.

We choose to model the probability of generating a *true or false belief* rather than the probability of generating a belief *for or against* ϕ because we have evidence for the former but not the latter based on the respective track-records of the two detectives. We assumed at the start of section 3 that this track-record information is known to the two detectives.

To avoid trivial cases, we assume that $0 < p < 1$ and $0 < q < 1$. We further assume that, if Jane or Hercule suspends judgment in response to

disagreement, their new opinion is generated with the same probabilities as their initial opinion (so Jane believes correctly with probability p , and Hercule believes correctly with probability q). We also assume that each time an opinion is generated this is done independently (in the probabilistic sense) from the detective's previous opinions and the other detective's current or previous opinions.

We think the assumption that the detectives reason independently from each other is justified because they make their assessments separately. If they are likely to come to the same conclusion this must be because the evidence points in a particular direction, which is reflected in the model by the choice of p and q .¹⁰ On the other hand, the assumption that newly generated opinions are independent from previously generated ones may be unrealistic, but it turns out not to have a strong influence on the results (see section 6 and appendix A).

In the epistemology of peer disagreement—as we learn from, for example, Christensen (2007, 216), Feldman (2007, 212), Elga (2007, 488), and Kelly (2010, 17)—the objective of rational conduct is commonly understood to be believing the correct truth-value. This suggests the following epistemic norm.

Accuracy Norm (AN). Having a true belief is more valuable than having a false belief.

We assume that Jane and Hercule share this noble goal, and that in fact obtaining a true belief about whether the butler did it is their *only* goal.¹¹ So the two detectives are not distracted by pragmatic concerns. This is a methodological rather than a substantive assumption: we are interested in the epistemology of peer disagreement, not its pragmatics.

¹⁰Note that p and q reflect the detectives' probabilities of reasoning correctly given the evidence, but not the probability of certain evidence being present; its presence or absence is taken as given for our purposes.

¹¹We recognize that one might have other epistemic goals than truth. In section 5 we show that some of these goals can be seen to follow from (AN). In section 6 we discuss the possibility of explicitly adding other norms.

(AN) determines the detectives' *preferences* over *outcomes* of the disagreement game: Jane prefers an outcome in which she has a true belief about the butler's guilt over one in which she has a false belief, and likewise for Hercule.¹² A detective receives utility 1 if her belief about the guilt or innocence of the butler at the end of the disagreement game is true, and utility 0 if it is false.¹³

The expected utility of a detective in the game is then simply the probability of ending the game with a true belief. So Jane and Hercule prefer a strategy if it increases their probability of ending the disagreement game with a true belief concerning ϕ .

We can now determine the probabilities of ending the disagreement game with a true belief for each combination of strategies of the two detectives (a combination of strategies is called a *strategy profile*).

If both detectives play **Stay**, they never change their mind in response to disagreement, so their probability of ending with a true belief is simply the probability that they obtain a true belief initially: p for Jane and q for Hercule. In all other cases the probability of ending the disagreement game with a true belief is the same for both detectives. These probabilities are indicated in table 1. The rows of table 1 indicate Jane's choice of strategy, and the columns indicate Hercule's choice.¹⁴

¹²Note that under our interpretation of (AN) detectives care only about the truth of their own belief. Results concerning a variation of our model where the detectives also care about the truth of the other detective's belief are available from the authors upon request.

¹³The introduction of utilities here adds nothing over and above the informal statement in the previous sentence. In particular the numbers 0 and 1 are arbitrary: all that matters is that a true belief yields a higher utility.

¹⁴This completes our specification of the game. Formally, a game is a triple $(N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$, where N is the set of players, S_i the set of strategies available to player i , and u_i the utility function for player i , which assigns real-valued utility to each strategy profile. In our case there are two players: $N = \{\text{Jane, Hercule}\}$; the strategy sets for both players are identical: $S_{\text{Jane}} = S_{\text{Hercule}} = \{\text{Stay, Suspend, Switch}\}$; and the utility for each player on each strategy profile is as in table 1.

The utilities are determined using the description of the disagreement game given in

| | Stay | Suspend | Switch |
|---------|----------|------------------------------------|------------------------------------|
| Stay | (p, q) | p | p |
| Suspend | q | $\frac{pq}{pq+(1-p)(1-q)}$ | $\frac{p(1-(1-p)(1-q))}{1-p(1-p)}$ |
| Switch | q | $\frac{q(1-(1-p)(1-q))}{1-q(1-q)}$ | pq |

Table 1: Expected utilities associated with each strategy profile under (AN).

How can the detectives maximize their probability of ending the disagreement game with a true belief, given that the choice of strategy of the other detective influences their probability of attaining true belief, but they cannot control it? Game theorists have invented various concepts of rationality in a game to deal with this problem. We will use the notion of Nash equilibrium.

A *Nash equilibrium* is a profile—that is, an assignment of a strategy to each player—in which either player’s strategy is a best response to the other’s. In other words, in a Nash equilibrium, no player can get an outcome she prefers over the equilibrium outcome by unilaterally changing her strategy. In our game this means that in a Nash equilibrium Jane and Hercule are maximizing their respective probabilities of ending the game with a true belief, *given* (that is, keeping fixed) the other detective’s strategy. This is how we interpret (epistemic) rationality for Jane and Hercule.

section 3. For example, if both detectives play **Suspend** they will generate new beliefs repeatedly until the first time they agree. The probability that they both generate a belief that ϕ is true is pq and the probability that they agree that ϕ is false is $(1-p)(1-q)$. So the probability that they end the game with a correct belief about ϕ is the probability that, on the first round on which they agree, they agree that ϕ is true rather than that ϕ is false. This probability is simply pq divided by $pq + (1-p)(1-q)$. See also appendix A.

5 Results and Discussion

What are the Nash equilibria of this game?¹⁵ This turns out to depend on the values of p and q . Figure 1 shows which strategy profiles are Nash equilibria for any combination of values of p and q .

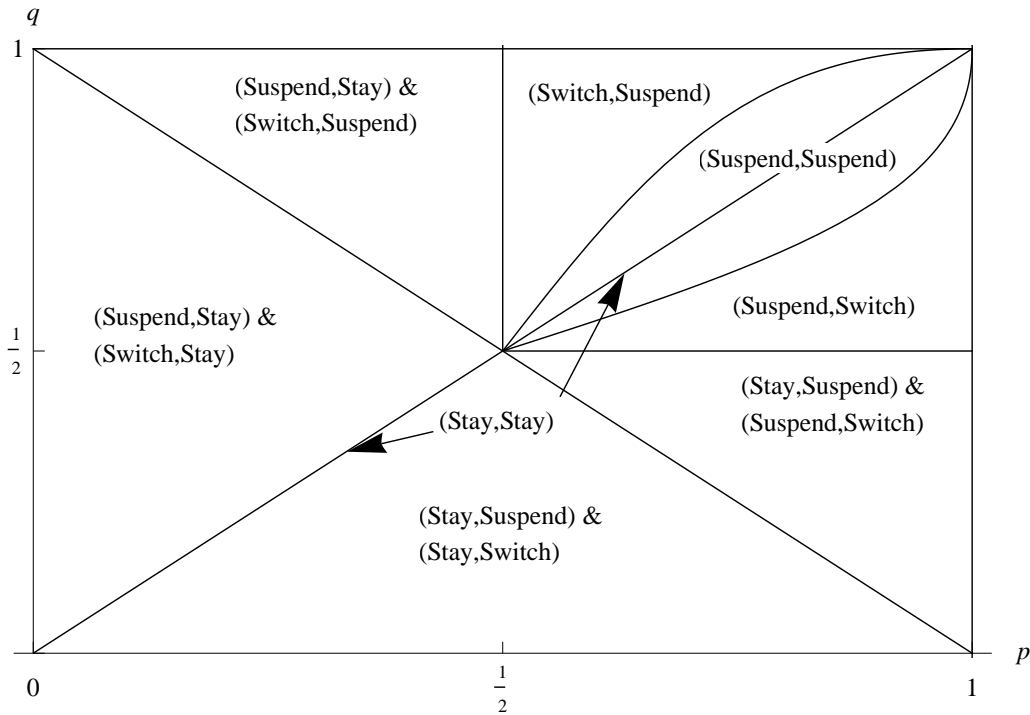


Figure 1: Nash equilibria of the peer disagreement game as a function of the truth-sensitivity of the detectives. E.g., if $p = 0.4$ and $q = 0.9$ then the Nash equilibria of the game are **(Suspend, Stay)** and **(Switch, Suspend)**.

Recall that we noted in section 3 that two factors would influence which strategy choice is best. First, the truth-sensitivity of the two detectives (modeled as p and q) and second, the strategy of the other detective. Both of these factors are shown in our results in figure 1.

¹⁵We consider only pure strategy equilibria.

The truth-sensitivity of the detectives clearly influences which strategy profiles are rational. For example, (**Stay,Switch**) is a Nash equilibrium whenever Hercule’s truth-sensitivity (his probability of drawing a true conclusion) is less than Jane’s truth-sensitivity and less than Jane’s probability of drawing a false conclusion (formally, $q \leq \min\{p, 1-p\}$). Similarly, (**Switch,Stay**) is a Nash equilibrium whenever Hercule’s truth-sensitivity is between Jane’s truth-sensitivity and Jane’s probability of drawing a false conclusion (formally, $p \leq q \leq 1-p$).

The other detective’s strategy also influences what it is rational for a detective to do. For example, if Hercule’s truth-sensitivity is higher than Jane’s probability of drawing a false conclusion, but less than $1/2$ (formally, $1-p \leq q \leq 1/2$), the Nash equilibria are (**Stay,Suspend**) and (**Suspend,Switch**). So under these circumstances, if Hercule chooses the strategy **Suspend**, it is rational for Jane to choose **Stay**, while if Hercule chooses **Switch**, it is rational for Jane to choose **Suspend**.

The epistemic success of the two detectives (both in terms of which strategy promises the best probability of a true belief, and in terms of the value of that probability) thus depends on the choices made by the other detective. In this way the epistemology of this model is truly *social*.

One way to understand the results in figure 1 is to view the detectives as making a tradeoff between two competing risks. On the one hand, there is the ‘cost’ of giving up one’s initial opinion. On the other hand, there is the cost of ignoring the other detective. When one detective has a significantly better track-record than the other (as reflected in the values of p and q), it is too costly for that detective to give up her initial opinion and switch to the other’s opinion. She gains more by staying with the initial belief, or suspending judgment and acquiring a new belief.

For the other detective it is the other way around. In her case, it is too costly to ignore the opinion of the other detective. Since she does not have as good a track-record, she would not gain as much by staying with her initial belief, or suspending judgment and acquiring a new belief, as she will by

switching to the opinion of the other detective. For her the cost of ignoring the other detective is higher than the cost of giving up her original opinion. The tipping points in these game-theoretic transactions can be read off from figure 1.

It is worth pointing out that the detectives' desire to minimize these risks is not epistemically basic. We have assumed that the only thing the detectives (ultimately) care about is maximizing their probability of ending the peer disagreement game with a true belief about ϕ . We now see that this goal, as formalized in (AN), implies that the detectives should worry about these two risks, and gives the detectives an epistemically motivated basis for trading them off against one another. In this sense our results fit nicely with the emerging literature that aims to explain various epistemic norms as following from (AN) (Joyce 1998, Pettigrew 2013).

Of particular interest in evaluating the results in figure 1 are the profiles (**Stay,Stay**) and (**Suspend,Suspend**). This is because the former captures most directly the Steadfast View—according to which it can be rational to **Stay** in a case of peer disagreement—and the latter captures most directly the Conciliatory View—according to which the only rational option is to **Suspend**.

What is surprising, and running contra the peer disagreement literature, is that *both* (**Stay,Stay**) and (**Suspend,Suspend**) turn out to constitute Nash equilibria, under some conditions even both at once.

As we can see from figure 1, the Steadfast profile (**Stay,Stay**) is a Nash equilibrium when Jane and Hercule are each other's equals in terms of how truth-sensitive their beliefs are (i.e., $p = q$). In such a case neither would gain anything by playing **Suspend** or **Switch** (provided the other detective continues to play **Stay**). This is because the probability that a detective ends up with a true belief by staying with her initial opinion is just as high as the probability that the opinion of the other detective or a newly generated opinion is true.

However, a mutual Conciliatory approach, as expressed in the strategy

profile (Suspend,Suspend), can *also* be a Nash equilibrium. This happens when p and q are both greater than $1/2$ and are relatively close to each other (see figure 1).¹⁶ When both detectives have relatively good track-records, and they find out that they have formed conflicting beliefs, they stand to gain more when they both suspend judgment and acquire a new belief, than when they stick to their initial beliefs, or switch to the other detective’s belief.

An especially interesting scenario occurs whenever p and q are exactly equal and greater than $1/2$: then (Stay,Stay) and (Suspend,Suspend) are Nash equilibria at the same time. Under the definition of rationality we use, in such a case both Steadfast and Conciliatory strategies are rational.

We wish to stress the significance of this result. In the literature on peer disagreement, the Steadfast strategy and the Conciliatory strategy are typically presented as mutually exclusive; *either* it is rational to play Stay *or* it is rational to play Suspend, but they cannot *both* be rational. A surprising insight of our analysis is that this need not be accurate. Under certain conditions, namely when two agents are positively and equally reliable, both the Steadfast strategy and the Conciliatory strategy can be rational. Moreover, the case where the two agents are positively and equally reliable is exactly the case the peer disagreement literature has focused on.

So where does this leave us in the peer disagreement debate? If we take seriously the modalities in the definitions of the views, the Steadfast View ‘wins’: it *can* be rational to stick to one’s opinion in the face of peer disagreement; the Conciliatory View’s claim that this cannot be rational is false in our model. But if we take the views as *recommending* strategies (Stay for the Steadfast View and Suspend for the Conciliatory View) then we think the Conciliatory View has the advantage—even when both are Nash equilibria—for the following reasons.

First, whenever (Stay,Stay) and (Suspend,Suspend) are Nash equilibria

¹⁶More precisely, the region where (Suspend,Suspend) is a Nash equilibrium is characterized by the inequality $\frac{p-\sqrt{p(1-p)}}{2p-1} \leq q \leq \frac{p^2}{1-2p(1-p)}$ (although the first expression is undefined when $p = 1/2$, the point $p = q = 1/2$ is also part of this region).

simultaneously, (**Suspend,Suspend**) offers a higher utility (a higher probability of solving the case correctly) to both detectives.¹⁷ In fact, (**Suspend,Suspend**) is *Pareto efficient*. So Jane and Hercule prefer to play (**Suspend,Suspend**) over (**Stay,Stay**). If they are allowed to discuss their strategy before the game starts, we should expect both detectives to play **Suspend**.

Second, **Suspend** is a *weakly dominant* strategy (for both detectives), while **Stay** is not. This means that playing **Suspend** pays off at least as well as playing **Stay** or **Switch**, *regardless* of what strategy the other detective chooses. So in this situation, playing **Stay** is only best for a detective who is absolutely certain that the other detective is playing **Stay** as well (and even then playing **Suspend** is equally good), whereas if there is only the slightest uncertainty about what the other detective is going to do, **Suspend** is the uniquely best strategy.

Third, we can see in figure 1 that when p and q are both greater than $1/2$ there is a significant area in which the profile (**Suspend,Suspend**) is a Nash equilibrium, while (**Stay,Stay**) is a Nash equilibrium only when p and q are exactly equal.¹⁸ This means that the strategy **Suspend** has a larger margin for error than the strategy **Stay**. If Jane and Hercule lack precise information about each other's truth-sensitivity (as is reasonable to expect), playing **Stay** is 'riskier' than playing **Suspend** because the former requires exact and the latter only approximate equality of the detectives' truth-sensitivities.

To sum up, a surprising result of this model is that if the detectives have equal track-records, and these track-records are 'good' (better than chance), then both the Steadfast profile and the Conciliatory profile are Nash equilibria. However, we have noted three reasons to think that in such cases the Conciliatory strategy should be preferred.

¹⁷Whenever $p = q > 1/2$, it must also be the case that $\frac{p^2}{p^2+(1-p)^2} > p$.

¹⁸More formally, the area where (**Stay,Stay**) is an equilibrium is measure zero in the parameter space, whereas the area where (**Suspend,Suspend**) is an equilibrium has positive measure.

6 Limitations and Extensions of Our Analysis

We have limited our analysis to a particular game-theoretic formalization of a particular disagreement game between two detectives, Jane and Hercule. To what extent does our analysis generalize to other peer disagreements? And what variations or extensions of our formalization are possible?

Regarding the first question, our analysis applies to peer disagreements in general insofar as they satisfy the assumptions of our model. In particular, (1) peers are cashed out in terms of comparable reliability or truth-sensitivity, (2) the possible responses available to the peers are something like the strategies **Stay**, **Suspend**, and **Switch** as we model them, and (3) the rationality of a particular response is evaluated in terms of how well it tracks the truth.

Regarding the second question, there are many options for different peer disagreement games. Let us give eight variables that can be filled in differently.

First, doxastic attitudes: in our model, strategies act on full belief states, but strategies might also be interpreted as adjusting degrees of belief.

Second, we forced our detectives to generate a new belief whenever they suspend judgment on ϕ . A variation of our model might allow peers to persist in a state of suspension. This outcome could be assigned its own value, presumably worse than having a true belief but better than having a false belief. We consider this variation in appendix B. Unsurprisingly, the results depend on what epistemic value is assigned to the state of suspension.

Third, we assumed that whenever Jane or Hercule generates a new belief (i.e., at the end of a round on which they disagreed and the relevant detective is playing **Suspend**) the new belief generated is probabilistically independent of the belief held on the previous round. This may seem unrealistic. For example, Jane may generally be a reliable detective ($p > 1/2$) but she may be prone to repeat mistakes in her reasoning. In appendix A we consider a version of the model in which newly generated beliefs are positively correlated

with the belief held on the previous round. The results are qualitatively similar to those of section 5.

Fourth, the number of peers. What happens if there are more than two disagreeing peers? Consider the case that we focused on above, where the peers' truth-sensitivity is equal, and greater than $1/2$. The Condorcet Jury Theorem shows that if a moderately large number of peers simultaneously state their opinion, the majority opinion is highly likely to be correct.¹⁹ But models of informational cascades show that if the peers state their opinion sequentially, the majority outcome is not nearly so informative (Bikhchandani et al. 1992, 996–999). This illustrates once again that the success of epistemic strategies—here majority voting, a plausible generalization of the Conciliatory View—can be quite sensitive to subtle contextual details, which formal models can focus attention on.

Fifth, we kept the peers' strategies fixed throughout the game. The reason for this was to enable an evaluation of the Conciliatory and Steadfast strategies. But it would be an interesting extension of the game to allow peers to change their strategies during the game.

Sixth, we assumed that the game might go on indefinitely. This is not very realistic. In real life there are time and energy constraints. So another possible extension would be to let the game continue for a limited number of rounds, after which the agents must have made up their minds. We consider the case with only one round in appendix B. If the other assumptions are unchanged, the results favor the Conciliatory View slightly more than those of the main text (see figure 3 in appendix B).

Seventh, in our analysis the rationality of a strategy was evaluated using Nash equilibria. Although this is very natural in game theory, it has substantive normative implications. So one may want to consider alternatives. Available alternatives include various refinements of the notion of equilibrium, such as the trembling hand equilibrium, and alternative standards,

¹⁹See List and Goodin (2001) for philosophical discussion and generalizations of the theorem.

such as weak dominance. Different strategies may turn out to be rational under such different standards of rationality.

Finally, we worked with only one epistemic norm, namely accuracy. But there are more epistemic goals. For example, many philosophers of science have argued, under the label of ‘epistemic diversity’, that maintaining diversity of opinion can have epistemic value to a population of scientists, stimulating new ideas and discoveries (Feyerabend 1975, Kitcher 1990, Zollman 2010). And the literature on epistemic rationality has identified a trade-off between truth and information (Levi 1967). For example, true beliefs could be maximized by believing only tautologies, but this is not informative. Either of these considerations could motivate augmenting or replacing (AN) with different norms.

7 Conclusion

By way of conclusion we emphasize four lessons that can be drawn from our preliminary game-theoretic investigation of the epistemology of peer disagreement.

First, in our model the Steadfast and Conciliatory strategies were sometimes both right: there were circumstances in which both staying with your own opinion and suspending belief were rational. The idea that staying and suspending can be rational simultaneously is underexplored in the literature and worth investigating more extensively.

Second, the rationality of a response to peer disagreement may depend on the truth-sensitivity of the peers. Both the peers’ relative truth-sensitivity (who has a better track-record and by how much?) and their absolute truth-sensitivity (are they better than chance, say, or some other objective threshold?) can make a difference.

Third, what is rational for a peer to do (e.g., whether to be Steadfast or Conciliatory) may depend on what the other peer is doing. This is a natural conclusion to draw in the game-theoretic context, but underexplored in the

peer disagreement literature.

Fourth, analysis of other game-theoretic models of peer disagreement may shed more light on the above three points and other important questions about peer disagreement. We encourage anyone interested in our model (especially if they liked it but for one or two assumptions) to develop and analyze such an alternative game-theoretic model of peer disagreement. We hope to have provided a fruitful framework within which such further models can be developed.

A A Model With Correlated Beliefs

In the main text we made a number of assumptions, some of which it might be desirable to relax or drop. In these appendices we consider a few slightly different versions of the model. In appendix A we relax the assumption that, when a detective plays **Suspend**, the new beliefs generated when a disagreement occurs are probabilistically independent of those generated on the previous round. In appendix B, we consider a non-iterative version of the game (i.e., the detectives' beliefs are evaluated at the end of the first round, rather than after a potentially lengthy exchange).

A.1 The Model

In the main text we assumed that whenever Jane or Hercule generates a new belief (i.e., at the end of a round on which they disagreed and the relevant detective is playing **Suspend**) the new belief generated is probabilistically independent of the belief held on the previous round. This may seem unrealistic. For example, Jane may generally be a reliable detective ($p > 1/2$) but she may be prone to repeat mistakes in her reasoning.

Here is a way to generalize the model to account for this possibility, while retaining the detectives' truth-sensitivity as a parameter of the model. Whenever a new belief needs to be generated (on a round other than the first), Jane retains the belief she had on the previous round with probability $1 - \delta$,

and she only generates a new belief with probability δ . If she generates a new belief, the new belief is true with probability p and false with probability $1 - p$.²⁰

As a result, Jane’s newly generated opinion is *positively correlated* with her opinion on the previous round, with lower values of δ corresponding to higher degrees of correlation. Similarly for Hercule: he keeps the belief he had on the previous round with probability $1 - \varepsilon$, and if he does generate a new belief it is true with probability q .

When $\delta = \varepsilon = 1$ this model reduces to the model of the main text in which newly generated beliefs are probabilistically independent of previous ones. Lower values of δ and ε correspond to higher degrees of correlation between old and new beliefs. Any level of (positive) correlation can be generated within this model (although we exclude the trivial case of perfect correlation that occurs when δ or ε is zero from the analysis). We do not consider anti-correlation.

What are the detectives’ payoffs, i.e., the probabilities of ending the game with a true belief? When neither detective plays **Suspend** nothing is really changed, so the payoffs are as in the original version of the peer disagreement game. The profiles (**Stay,Suspend**) and (**Suspend,Stay**) are also unchanged: as long as there is some probability of the detectives changing their mind when playing **Suspend** (i.e., as long as $\delta > 0$ and $\varepsilon > 0$), in these profiles the detectives eventually agree on whichever opinion the detective playing **Stay** started with (with probability one).

For the three remaining profiles, determining the payoffs is a little more involved. Here we rely on an analysis using (absorbing) *Markov Chains*. On any round, the game may be in one of four states: Jane has a true belief and Hercule has a false belief (state s_1), Jane has a false belief and Hercule has a true belief (s_2), Jane and Hercule both have a true belief (s_3), or Jane

²⁰Of course, only one of these two latter possibilities will result in a genuinely new belief: even in the case where she generates a “new” belief there is some chance that this belief happens to be the same as the one she had on the previous round.

and Hercule both have false beliefs (s_4). States s_1 and s_2 are *transient*, while state s_3 and s_4 are *absorbing*: the game ends when an absorbing state is reached. The probabilities that the game is in one of these states in the first round are given by the (row) vectors \mathbf{v}_d and \mathbf{v}_a , where

$$\mathbf{v}_d = \begin{matrix} & s_1 & s_2 \\ \left[\right. & p(1-q) & (1-p)q \end{matrix} \quad \text{and} \quad \mathbf{v}_a = \begin{matrix} & s_3 & s_4 \\ \left[\right. & pq & (1-p)(1-q) \end{matrix}.$$

Following standard practice for absorbing Markov Chains, we write the matrix of transition probabilities P as a block matrix:

$$P = \begin{bmatrix} Q & R \\ O & I \end{bmatrix}.$$

Here, Q is the matrix of transition probabilities between the transient states, R is the matrix of transition probabilities from the transient states to the absorbing states, O is the 2×2 zero matrix, and I is the 2×2 identity matrix (because there are no transitions out of absorbing states).

If both detectives play **Suspend**, the transition probabilities from the transient states are as follows:

$$Q = \begin{matrix} & s_1 & s_2 \\ \begin{matrix} s_1 \\ s_2 \end{matrix} \left[\right. & (1 - \delta(1-p))(1 - \varepsilon q) & \delta(1-p)\varepsilon q \\ & \delta p \varepsilon(1-q) & (1 - \delta p)(1 - \varepsilon(1-q)) \end{matrix},$$

$$R = \begin{matrix} & s_3 & s_4 \\ \begin{matrix} s_1 \\ s_2 \end{matrix} \left[\right. & (1 - \delta(1-p))\varepsilon q & \delta(1-p)(1 - \varepsilon q) \\ & \delta p(1 - \varepsilon(1-q)) & (1 - \delta p)\varepsilon(1-q) \end{matrix}.$$

For example, the transition probability from s_1 to s_2 is the probability that Jane moves from a true to a false belief, which is $\delta(1-p)$, times the probability that Hercule moves from a false to a true belief, which is εq .

Now the probabilities of ending the game in each of the absorbing states s_3 and s_4 are given in a row vector \mathbf{w} , where

$$\mathbf{w} = \mathbf{v}_d \left(\sum_{k=0}^{\infty} Q^k \right) R + \mathbf{v}_a = \mathbf{v}_d (I - Q)^{-1} R + \mathbf{v}_a.$$

The first entry of \mathbf{w} is the probability of ending the game in the absorbing state where both detectives have a true belief, and the second entry is the probability of both detectives ending up with a false belief.²¹ Hence the (expected) payoff to Jane and Hercule on the profile (**Suspend,Suspend**) is

$$\mathbf{w}_1 = \frac{pq(\delta\varepsilon + \delta^2(1-\varepsilon)(1-p) + (1-\delta)\varepsilon^2(1-q))}{\delta\varepsilon(pq + (1-p)(1-q)) + \delta^2(1-\varepsilon)p(1-p) + (1-\delta)\varepsilon^2q(1-q)}. \quad (1)$$

We can perform the same analysis for the profile (**Suspend,Switch**). The only difference is in the matrices of transition probabilities Q and R . When Hercule plays **Switch**, he switches back and forth between a true and a false belief on any round in which the game remains in a transient state (once the detectives agree—i.e., an absorbing state is reached—he stops switching). Hence

$$Q = \begin{matrix} & \begin{matrix} s_1 & s_2 \end{matrix} \\ \begin{matrix} s_1 \\ s_2 \end{matrix} & \begin{bmatrix} 0 & \delta(1-p) \\ \delta p & 0 \end{bmatrix} \end{matrix} \quad \text{and} \quad R = \begin{matrix} & \begin{matrix} s_3 & s_4 \end{matrix} \\ \begin{matrix} s_1 \\ s_2 \end{matrix} & \begin{bmatrix} 1 - \delta(1-p) & 0 \\ 0 & 1 - \delta p \end{bmatrix} \end{matrix}.$$

Using the same formula for \mathbf{w} , we find that the payoff to Jane and Hercule on the profile (**Suspend,Switch**) is

$$\mathbf{w}_1 = p \frac{1 + \delta(1-p)(2q - \delta q - 1)}{1 - \delta p(1-p)}. \quad (2)$$

By symmetry, the payoff on the profile (**Switch,Suspend**) is

$$q \frac{1 + \varepsilon(2p - \varepsilon p - 1)(1-q)}{1 - \varepsilon q(1-q)}. \quad (3)$$

²¹When at least one detective plays **Suspend** and $\delta > 0$ and $\varepsilon > 0$, the detectives eventually end up agreeing with probability one. As a result, the two entries of \mathbf{w} sum to one.

Hence the full table of expected utilities is as given in table 2. Note that if $\delta = \varepsilon = 1$ we are back in the case we started with, where new beliefs are generated independently of those generated in previous rounds. The above derivation thus also serves as a proof of the values given in table 1.

| | Stay | Suspend | Switch |
|---------|----------|---------|--------|
| Stay | (p, q) | p | p |
| Suspend | q | (1) | (2) |
| Switch | q | (3) | pq |

Table 2: Expected utilities in the peer disagreement game with correlated beliefs.

A.2 Results and Discussion

What are the Nash equilibria of this generalized version of the peer disagreement game? This depends on the values of p and q (as before) but also on the values of δ and ε . Suppose, for instance, that $\delta = \varepsilon = 1/4$, i.e., in case the two detectives disagree and one of them plays **Suspend**, she copies her reasoning from the previous round three quarters of the time, generating a new belief in accordance with her truth-sensitivity only one quarter of the time. The Nash equilibria of this game are as shown in figure 2.

The results are qualitatively quite similar to those in the original peer disagreement game. In particular, **(Stay,Stay)** is an equilibrium if and only if $p = q$, and **(Suspend,Suspend)** is an equilibrium in the top-right corner of the figure. The case where **(Stay,Stay)** and **(Suspend,Suspend)** are equilibria simultaneously occurs whenever $p = q \geq 0.57$. Figure 2 only shows the case where $\delta = \varepsilon = 1/4$, but to a large extent these features generalize to all possible values of δ and ε , as the following propositions show.

Proposition 1. *In the peer disagreement game with correlated beliefs, **(Stay,Stay)** is a Nash equilibrium if and only if $p = q$.*




Figure 2: Nash equilibria of the peer disagreement game with correlated belief ($\delta = \varepsilon = 1/4$) as a function of p and q .

Proposition 2. *In the peer disagreement game with correlated beliefs, if $p = q \geq 0.63$ and $\delta = \varepsilon$ then $(\text{Suspend}, \text{Suspend})$ is a Nash equilibrium.*

Recall from figure 1 that at least one of $(\text{Suspend}, \text{Suspend})$, $(\text{Suspend}, \text{Switch})$, and $(\text{Switch}, \text{Suspend})$ was a Nash equilibrium if and only if $p + q \geq 1$. Arguably, all three of these profiles reflect a version of the Conciliatory View (after all, they call on both detectives to revise their view in light of the disagreement). While the region where $(\text{Suspend}, \text{Suspend})$ is an equilibrium is quite sensitive to variations in the values of δ and ε , this is only due to competition from these other two “Conciliatory” profiles. The region where at least one of these three Conciliatory profiles is a Nash equilibrium turns out to be robust to the introduction of correlated beliefs.

Proposition 3. *In the peer disagreement game with correlated belief, if $p + q \geq 1$ then at least one of the Conciliatory profiles is a Nash equilibrium.*

Proposition 4. *In the peer disagreement game with correlated belief, if $p = q > 1/2$ then both detectives have a higher probability of ending the game with a true belief under any of the three Conciliatory profiles than under (Stay,Stay).*

So in the game with correlated beliefs we can reproduce all the important features of the game in the main text: (1) (Stay,Stay) is only an equilibrium if $p = q$, (2) Conciliatory and Steadfast profiles can be Nash equilibria simultaneously, (3) Conciliatory equilibria are more robust to uncertainty about truth-sensitivity, (4) when there are both Conciliatory and Steadfast equilibria, Conciliatory equilibria are preferable by the detectives' own lights.

We conclude that introducing correlations between the detectives' beliefs on different rounds does not have a strong influence on the results. In particular, all of the qualitative conclusions we drew about the peer disagreement game in section 5 remain valid when such correlations are introduced.

B A Model With Only One Round

B.1 The Model

Another potentially problematic assumption of our model is its iterative nature. In particular, objections have been raised against the assumption that the strategy **Suspend** involves, among other things, gathering new evidence (reflected by a new independent draw from the probability distribution by which initial beliefs were generated). This may be thought of as being against the spirit of the peer disagreement debate.

In this appendix we will consider two versions of the model that respectively weaken and completely remove that assumption. The first model moves from an iterative model to a model with only one round, but keeps the assumption that if the detectives **Suspend** on that one round, they gather new

evidence to form their final belief. The second model also has only one round, but now the gathering of new evidence is removed in favor of the possibility of ending the game in a suspended state.

How are the probabilities of ending the game with a true belief affected by changing the model to have only one round? Details are given in table 3. Since in a one round game only the detective's own strategy matters, this table just gives Jane and Hercule's payoff as a function of their own strategy.

| | Stay | Suspend | Switch |
|---------|------|-------------------------------|--------|
| Jane | p | $pq + p(p(1 - q) + (1 - p)q)$ | q |
| Hercule | q | $pq + q(p(1 - q) + (1 - p)q)$ | p |

Table 3: Expected utilities associated with each strategy if there is only one round.

The probabilities in table 3 are obtained as follows. If the detective plays **Stay**, the probability of ending the game with a true belief is the probability of generating a true belief at the start of the first round (p for Jane, q for Hercule). If the detective plays **Switch**, at the end of the first round she always has the belief the other detective started with, so the probability of ending the game with a true belief is the probability that the other detective got it right initially (q for Jane, p for Hercule). If the detective plays **Suspend**, there are two ways to end the game with a true belief. Either they immediately agree on the correct truth-value (this happens with probability pq) or they disagree (probability $p(1 - q) + (1 - p)q$) but the detective suspends and the new belief she generates is correct.

Now consider a model in which, as above, there is only one round of the game, but if a detective suspends she does not generate a new belief. The question is how ending the game in a state of suspension should be evaluated. We introduce a new parameter u_s to capture this value. It seems clear that being in suspension is worse than having a true belief but better than having a false belief, which implies that $0 < u_s < 1$. The most reasonable value of

u_s is perhaps $1/2$ but we will not assume this.

| | Stay | Suspend | Switch |
|---------|------|---------------------------------|--------|
| Jane | p | $pq + u_s(p(1 - q) + (1 - p)q)$ | q |
| Hercule | q | $pq + u_s(p(1 - q) + (1 - p)q)$ | p |

Table 4: Expected utilities associated with each strategy if no new evidence is gathered.

Payoffs to the two detectives are given in table 4. The only change compared to table 3 is that if a detective plays **Suspend** and the detectives disagree (with probability $p(1 - q) + (1 - p)q$) she gets a payoff of u_s (the value of ending the game in suspension) rather than p or q (the probability that she would obtain a true belief if allowed to gather new evidence).

B.2 Results and Discussion

What are the Nash equilibria of the version of the peer disagreement game with only one round whose payoffs are given in table 3, i.e., the version in which the detectives gather new evidence if they suspend judgment (so the value of suspending is irrelevant)? As before, it depends on the values of p and q , as shown in figure 3.

The results are similar to those obtained before. The two main differences are that (1) **Suspend** is no longer a good strategy outside of the top-right corner and that (2) the profile (**Stay,Stay**) no longer appears as a Nash equilibrium in the top-right corner, even when $p = q$.

Particularly noteworthy is that (**Suspend,Suspend**) remains a Nash equilibrium in the situations most directly relevant to the peer disagreement debate: when p and q are both greater than $1/2$ and relatively close to each other. Hence we take our results here to favor the Conciliatory View slightly more than the results obtained in the main text.

What if no new evidence is gathered, i.e., payoffs are as in table 4? Then




Figure 3: Nash equilibria of the peer disagreement game with only one round.

the Nash equilibria of the game depend on the value u_s of being in a state of suspension, in addition to the values of p and q . This is illustrated for a number of values of u_s in figure 4.

If $u_s < 1/2$, it is never a good idea to suspend judgment: at least one of the alternative strategies **Stay** and **Switch** always leads to a better payoff. If $u_s = 1/2$, the three strategies **Stay**, **Suspend**, and **Switch** pay off equally well whenever $p = q$, so in that case any of the nine strategy profiles, including **(Suspend,Suspend)**, is a Nash equilibrium on the line through the middle of figure 4 (top-left). If $u_s > 1/2$, **Suspend** is the unique best strategy whenever p and q are relatively close to each other in value (where what it means to be “close” is stricter when u_s is closer to $1/2$ and looser when u_s is closer to 1). Accordingly, as u_s increases, we see a growing area in the middle of

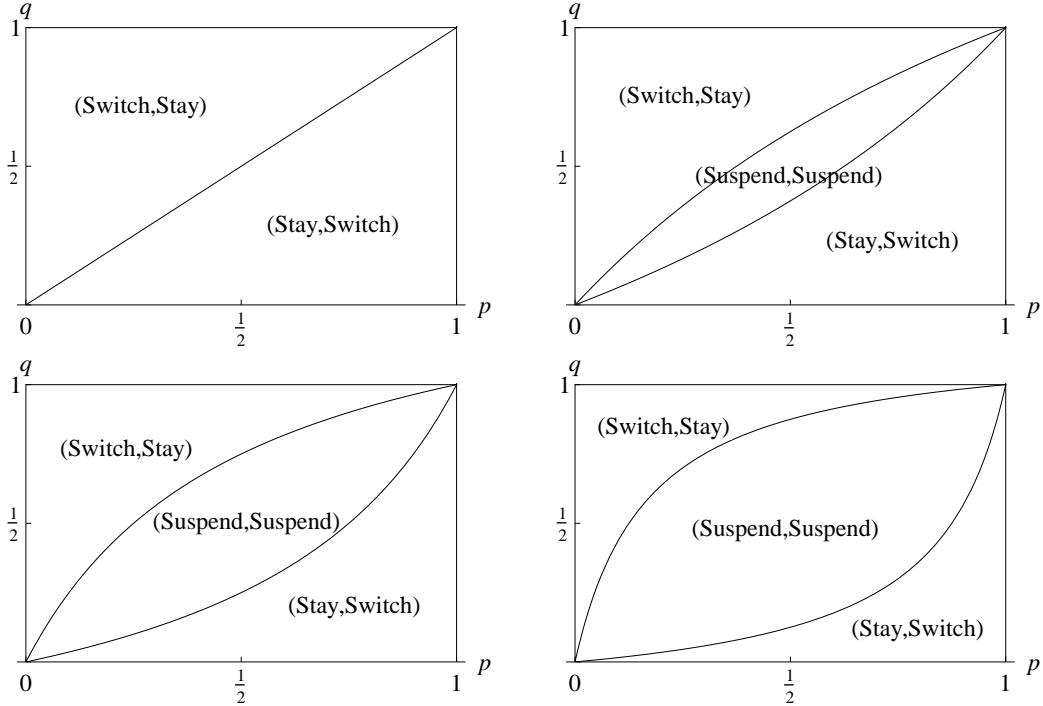


Figure 4: Nash equilibria of the peer disagreement game when $u_s \leq \frac{1}{2}$ (top-left), when $u_s = \frac{5}{8}$ (top-right), when $u_s = \frac{3}{4}$ (bottom-left), and when $u_s = \frac{7}{8}$ (bottom-right).

the figures where **(Suspend,Suspend)** is the unique Nash equilibrium.

Unlike either the original iterative model or the version of the one round model considered above, the area where **(Suspend,Suspend)** is an equilibrium is not confined to the upper right corner of the graph.

How do these results relate to the peer disagreement debate? Regarding the first one-round model, whereby the agents are allowed to re-examine the evidence, we can see that the Steadfast strategy becomes less rewarding. The strategy **Stay** no longer appears in the top-right corner of figure 3, which is the area that is interesting for the debate about peer disagreement. This means that, when the detectives are peers in that they have comparably good track-records, the Steadfast strategy **Stay** is always outclassed by other

strategies, most importantly by the Conciliatory profile (**Suspend,Suspend**).

Next, consider the second one-round model, whereby the agents do not go back to the evidence to form a new opinion. Again we look at the more interesting top-right corners in figure 4. Somewhat unsurprisingly, we can see that which strategy is best depends on the value that is assigned to suspension of judgment. The more valuable this is considered to be, the more rewarding a Conciliatory strategy becomes, and the less valuable suspending is, the more a Steadfast strategy comes into play again.

References

- Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50(2):510–530, 1985.
- Robert J. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.
- Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, 1992.
- Gustavo Cevolani. Truth approximation, belief merging, and peer disagreement. *Synthese*, 191(11):2383–2401, 2014.
- David Christensen. Epistemology of disagreement: The good news. *The Philosophical Review*, 116(2):187–217, 2007.
- David Christensen. Disagreement as evidence: The epistemology of controversy. *Philosophy Compass*, 4(5):756–767, 2009.
- Morris H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

- Kenny Easwaran. Dr. Truthlove or: How I learned to stop worrying and love Bayesian probabilities. *Noûs*, forthcoming. doi: 10.1111/nous.12099.
- Adam Elga. Reflection and disagreement. *Noûs*, 41(3):478–502, 2007.
- Lee Elkin. An epistemically modest response to disagreement, AGM-ified. *The Reasoner*, 9(9):76–77, 2015.
- Richard Feldman. Reasonable religious disagreements. In Louise Antony, editor, *Philosophers Without Gods: Meditations on Atheism and the Secular*, pages 194–214. Oxford University Press, 2007.
- Paul Feyerabend. *Against Method*. New Left Books, London, 1975.
- John R. P. French, Jr. A formal theory of social power. *Psychological Review*, 63(3):181–194, 1956.
- John D. Geanakoplos and Heraklis M. Polemarchakis. We can’t disagree forever. *Journal of Economic Theory*, 28(1):192–200, 1982.
- Stephan Hartmann and Jan Sprenger. Judgment aggregation and the problem of tracking the truth. *Synthese*, 187(1):209–221, 2012.
- Stephan Hartmann, Gabriella Pigozzi, and Jan Sprenger. Reliable methods of judgement aggregation. *Journal of Logic and Computation*, 20(2):603–617, 2010.
- David Jehle and Branden Fitelson. What is the “Equal Weight View”? *Episteme*, 6(3):280–293, 2009.
- James M. Joyce. A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4):575–603, 1998.
- Thomas Kelly. The epistemic significance of disagreement. In Tamar Szabó Gendler and John Hawthorne, editors, *Oxford Studies in Epistemology*, volume 1, pages 167–196. Oxford University Press, 2005.

- Thomas Kelly. Peer disagreement and higher-order evidence. In Richard Feldman and Ted Warfield, editors, *Disagreement*, pages 111–174. Oxford University Press, 2010.
- Philip Kitcher. The division of cognitive labor. *The Journal of Philosophy*, 87(1):5–22, 1990.
- Jennifer Lackey. What should we do when we disagree? In Tamar Szabó Gendler and John Hawthorne, editors, *Oxford Studies in Epistemology*, volume 3, pages 274–93. Oxford University Press, 2008.
- Maria Lasonen-Aarnio. Disagreement and evidential attenuation. *Nous*, 47(4):767–794, 2013.
- Keith Lehrer and Carl Wagner. *Rational Consensus in Science and Society: A Philosophical and Mathematical Study*, volume 24 of *Philosophical Studies Series in Philosophy*. D. Reidel, Dordrecht, 1981.
- Hannes Leitgeb. The stability theory of belief. *Philosophical Review*, 123(2):131–171, 2014.
- Isaac Levi. *Gambling with Truth*. MIT Press, Cambridge, 1967.
- Hanti Lin and Kevin T. Kelly. Propositional reasoning that tracks probabilistic reasoning. *Journal of Philosophical Logic*, 41(6):957–981, 2012.
- Christian List. Social choice theory. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Winter 2013 edition, 2013. URL <http://plato.stanford.edu/archives/win2013/entries/social-choice/>.
- Christian List and Philip Pettit. Aggregating sets of judgments: An impossibility result. *Economics and Philosophy*, 18:89–110, 2002.
- Christin List and Robert E. Goodin. Epistemic democracy: Generalizing the Condorcet Jury Theorem. *Journal of Political Philosophy*, 9(3):277–306, 2001.

- Carlo Martini, Jan Sprenger, and Mark Colyvan. Resolving disagreement through mutual respect. *Erkenntnis*, 78(4):881–898, 2013.
- Sarah Moss. Scoring rules and epistemic compromise. *Mind*, 120(480):1053–1069, 2011.
- Richard Pettigrew. Epistemic utility and norms for credences. *Philosophy Compass*, 8(10):897–908, 2013.
- Jan-Willem Romeijn. Pooling, voting, and Bayesian updating. Manuscript, 2015.
- Peter van Inwagen. We’re right. they’re wrong. In Richard Feldman and Ted Warfield, editors, *Disagreement*, pages 10–28. Oxford University Press, 2010.
- Roger White. Epistemic permissiveness. *Philosophical Perspectives*, 19(1):445–459, 2005.
- Kevin J. S. Zollman. The epistemic benefit of transient diversity. *Erkenntnis*, 72(1):17–35, 2010.