

Conciliationism and Merely Possible Disagreement¹

Zach Barnett and Han Li

Abstract: Conciliationism faces a challenge that has not been satisfactorily addressed. There are clear cases of epistemically significant merely possible disagreement, but there are also clear cases where merely possible disagreement is epistemically irrelevant. Conciliationists have not yet accounted for this asymmetry. In this paper, we propose that the asymmetry can be explained by positing a selection constraint on all cases of peer disagreement—whether actual or merely possible. If a peer’s opinion was not selected in accordance with the proposed constraint, then it lacks epistemic significance. This allows us to distinguish the epistemically significant cases of merely possible disagreement from the insignificant ones.

1. Introduction

Conciliationism is a view about how to respond rationally to the discovery of disagreement. You believe some proposition P. You discover that someone else believes the opposite. Provided that certain constraints are met,² Conciliationism (as we’ll understand it) holds that you should suspend judgment about P.³ These constraints can be made precise in various ways, but for this paper, it won’t matter. No matter how the constraints are elaborated, Conciliationism faces a problem that has not been

¹ For comments on earlier drafts and helpful discussion of relevant issues, we would like to thank David Black, David Christensen, Harry Chalmers, Nina Emery, Georgi Gardner, Joshua Schechter, and audiences at the Princeton/Rutgers Graduate Conference 2014, the APA Eastern Division 2014, the Brown University Philosophy Graduate Forum, and a seminar on the etiology of belief which met at Brown University in the fall of 2013.

² The constraints typically involve a person’s having good reason to consider her disagreeer(s) at least as rational, or at least as reliable with respect to the disputed sort of issue as she is herself. Disagreeers (or agreeers) occupying this role are sometimes termed “epistemic peers,” though Conciliationists have pointed out that at least some reduction in confidence can be required even under less stringent conditions. For a sample of what some have said on this issue, see Feldman and Warfield (2009) and Christensen and Lackey (2013).

³ “Conciliationism” refers to a broad family of views, according to which, a person should, upon discovering disagreement of a certain sort, revise her opinion in the direction of her disagreeer. Some such views apply to credences, or partial beliefs, and typically require a person to reduce confidence in a disputed proposition. Other views apply to all-or-nothing belief, and often recommend suspension of judgment. We rely on the all-or-nothing conception in order to cast the central challenge resolved in this paper in the starkest possible terms.

satisfactorily addressed. Kelly raises the problem in his (2005) discussion of merely possible disagreement. Kelly writes:⁴

[W]hether there is any actual disagreement with respect to some question as opposed to merely possible disagreement might, in a particular case, be an extremely contingent and fragile matter. In particular, whether there is any actual disagreement might very well depend on factors that everyone will immediately recognize as irrelevant to the truth of the question at issue.

Kelly's thought is something like this: Some possible disagreements become actualized. Some don't. And it is something of an accident which ones become actual. Why should an accident have any effect on one's rational confidence? The below case makes the point vivid:⁵

Evil Tyrant: A powerful, prophetic tyrant, decides that no one should ever believe the proposition *P*, so he orders the execution of all would-be *P*-believers (all of the people who would come to believe *P*, if they were permitted to live). 1,000 people are executed. Some time later, there is something of a consensus about whether *P*. The 1,000 remaining people all believe $\sim P$. There *could have been* disagreement about *P* if the tyrant had acted differently, but in fact there is only agreement.

Even though there is no actual disagreement about *P*, one would expect Conciliationists to hold that survivors (who know exactly what the tyrant did) should be deeply worried about their belief in $\sim P$ in this case. Knowing that certain people would have disagreed, had they been left alive, seems to generate substantial conciliatory pressure. As apparent as it may be to the Conciliationist that survivors ought to suspend judgment in this case, justifying this intuition is not a trivial task. Certainly,

⁴ Kelly (2005) p. 18.

⁵ This case was originally suggested, in outline, by Kelly (2005, p. 18).

Conciliationists cannot affirm that merely possible disagreement has the same epistemic significance as actual disagreement in all cases. If they were to, they would be forced to suspend judgment far too often. Consider a more straightforward case:

Bare Possibility: You believe that P. You realize a peer could have believed \sim P.

Intuitively, suspension of judgment does not seem to be required here. Moreover, given our fallibility, it seems clear that for almost any proposition one believes, it will be true that a peer could believe its negation. If one were to suspend judgment in all such cases, there would be very few propositions about which one could rationally have any opinion — unacceptably few.

So the Conciliationist faces a difficult question: What distinguishes Evil Tyrant from Bare Possibility? Both are cases of possible disagreement; neither is a case of actual disagreement. If there is an important difference between these two cases, the Conciliationist should say something about what that difference is. Henceforth, we refer to this problem as the “Modal Challenge” for Conciliationism.⁶

As we see it, the Modal Challenge has not yet been answered. Admittedly, some Conciliationists have addressed the issue of merely possible disagreement in general. With few exceptions, they contend that merely possible disagreement carries little, if

⁶ What we refer to as the “Modal Challenge” should be distinguished from other challenges to Conciliationism involving merely possible disagreement from Kelly (2005). Specifically, Kelly discusses numerous examples involving Newcomb’s Problem that have received attention from Conciliationists. See, for example, Carey (2011), Christensen (2007), and Kornblith (2010). We will not be discussing these examples, as they raise distinct issues from the ones that arise here. The Modal Challenge deserves individual attention.

any epistemic significance.⁷ While this may be true in most cases, Evil Tyrant shows that it cannot always be true. Under what conditions does merely possible disagreement have epistemic significance?⁸ This paper proposes a solution to this problem. The proposal involves positing a selection constraint on all disagreements, which disqualifies cases like Bare Possibility but not cases like Evil Tyrant. Merely possible disagreement has epistemic significance when, and only when, it was *sampled* in accordance with certain statistical constraints. Importantly, the proposed constraints cut across the actual/possible distinction. We illustrate the idea through an analogy.

2. Biased Coins, Biased Thinkers

You are in a warehouse containing one million coins. All of the coins have a proposition P written on one side and its negation on the other. You start with one such coin in your hand. You know that all of the coins (including yours) are not fair but in fact biased toward the side with the true proposition on it. You could investigate whether P simply by flipping your coin and observing the result. Of course, this isn't a

⁷ See, for example, Christensen (2007), Carey (2013), Frances (2010).

⁸ Carey and Matheson (2013) provide the only discussion we are aware of which makes any direct contact with this question. According to their view, evidence of merely possible disagreement can have epistemic significance so long as the possible disagreement is "in some important sense, 'nearby.'" More specifically, they hold that if there are "many nearby worlds" in which the consensus opinion diverges from the actual consensus, then this provides reason to doubt the actual consensus view. So in essence, they think that merely possible disagreement is sometimes but not always significant. Ultimately, the proposal we recommend is consonant with the Carey and Matheson view in certain broad respects. But importantly, as we will see, a satisfying solution to the Modal Challenge can be given which relies primarily on what we know from statistics, and on little else. No grappling with contentious questions about 'nearbyness' or the metaphysics of possible worlds more generally is necessary to solve this problem.

perfect method. Although the coin has a higher chance of landing truly (i.e. landing on the side with the true proposition), it still might land falsely. Ordinarily, you might flip the coin many times in order to gain better and better evidence about whether P. But suppose that each coin can be flipped only once. In this situation, you could improve upon the evidence provided by the first flip by flipping the other coins. Since each coin flip is independent, you could eventually obtain very good evidence about whether P by flipping enough of them.

Now imagine a different case. Suppose P is a true proposition. As a somewhat reliable thinker, you are more likely to arrive at P than at \sim P. Like the coin, you are “biased” toward the truth.⁹ You have some reason to believe the result of your inquiry, but you recognize that you may be mistaken. And just as with the coin, you can’t perform additional “trials” on yourself – at least not in any straightforward manner. Now suppose that there are other thinkers who are also biased toward the truth. By surveying them, you can gain more evidence about whether P, and with enough of them, your evidence can be very good.¹⁰ This is one way to understand the epistemic significance of peer opinion.¹¹

⁹ Note that “biased” is being used in an unusual sense.

¹⁰ For simplicity, let’s assume for now that the thinkers, like the coins, are statistically independent of each other. We will return to this issue later.

¹¹ The foregoing passage seems to suggest that a person should view herself as merely another reliable indicator of truth among many. We opt to frame it this way solely for ease of exposition, officially staying neutral about whether a person’s own judgment screens her evidence. For discussion of this sticky issue, see White (2009), Enoch (2010), Sliwa and Horowitz (forthcoming), Aarnio (forthcoming), Weatherson (ms.),

Roughly speaking, what you are doing in both cases is taking a sample of similarly good indicators of the truth. As the sample size increases, the most common outcome is increasingly likely to correspond with the truth. Understanding disagreement in terms of sampling is helpful to fully answer our original question. Let's return to the coin warehouse in order to consider a parallel to Evil Tyrant:

Coin Tyrant: You flip the biased coin in your hand once and it lands P. You find other coins and you flip them. If Coin Tyrant realizes that a coin is going to land $\sim P$, he vaporizes it while it is in the air. This happens many times. You never observe any coin actually land $\sim P$.

No coins actually disagree with yours. But it is intuitively apparent that you should take those would-have-been flips into account when you interpret the data. Why? Because learning about how the vaporized coins would have landed seems to give just as much evidence about whether P as you would have had if they had actually been permitted to land. Now consider an intuitively uninformative case of merely possible coin disagreement:

Bare Possibility (Coin): You have flipped the coin in your hand once, and you obtained P. You find another coin in the warehouse. Before you flip it, you realize that it could land $\sim P$.

It is clear that you should be unmoved in this case. You should not suspend judgment about whether to believe P simply on the basis of its being possible that another similar coin could land differently than yours did. This possibility seems to tell you nothing about how that other coin is *likely* to land, which is, of course, what matters when we are trying to inquire about P with these coins.

So how do we explain our asymmetrical intuitions in these cases? As we said

above, we want to understand both cases in terms of statistical sampling. In both, we are trying to estimate a certain value – namely, P 's truth-value. We know that the coins (and the thinkers) are more likely to arrive at the right belief than not. As long as the trials are independent of each other, the outcome that occurs most often over some sample is more likely to correspond to the truth. And the bigger the sample, the better. But in order for an outcome to count as evidence in a sample, the selection procedure that was used must meet certain constraints. A sample exclusively of NBA players, for instance, would not provide good evidence about the average height of a US citizen. In short, our view is that Bare Possibility (Coin) and its disagreement analogue involve bad selection procedures of this sort while Coin Tyrant and its disagreement analogue don't. In the next section, we defend this analogy between coins and thinkers, which, as we have seen, suggests an answer to The Modal Challenge.

3. Defending the Analogy

The solution we favor ultimately depends upon the idea that thinkers are, in certain relevant respects, like coins. With the coins, it was very clear from which outcomes we were sampling. It was very clear that the coins had certain probabilities associated with these outcomes, and, moreover, that the probabilities could be arranged such that the coins were about equally biased toward the truth. It was very clear that the coins were statistically independent. And most importantly, as a result of these properties, it was clear that one could accurately estimate P 's truth-value by flipping

these coins and taking a sample. As we will now argue, something similar can be said for thinkers. Specifically, we hold that in paradigmatic cases of disagreement that interest Conciliationists, it makes sense to a person in such a case to view her fellow thinkers as similarly reliable, and (more or less) independent indicators of truth. And as a result, taking a sample of thinker's opinions can be a good way for her to investigate the truth of P.

To begin, which outcomes are we sampling from? It is easy enough to see the different outcomes from which we are sampling when we are working with coins: Heads and Tails. In the case of thinkers, it is natural to sample from the set of doxastic states that the thinker could take toward P. We will individuate doxastic states in terms of all-or-nothing beliefs.¹² Accordingly, we think of each situation as having three possible outcomes: belief toward P, disbelief toward P, and suspension of judgment toward P.

Next, let us discuss the probabilities associated with these outcomes. Why think that what a thinker will believe is a probabilistic matter at all? We can lean on common sense to answer this question. Specifically, we make everyday probability judgments about what one could have believed or would have been likely to have believed. For example, suppose I discover that I made a simple error on some math problem. I might be inclined to say something like "It was very unlikely that I would have made such an

¹² We suspect our view can be squared with credences as well. This involves making sense of credences about probabilities. See Lewis (1980) for a discussion on this subject.

egregious error. I usually am good at carrying my ones." Or we might say "It is unlikely that Nancy Drew would come to a conclusion about the murder given the sparseness of the evidence." When we make statements like these, we seem to be assuming that there was more than one belief a thinker might have arrived at and that some outcomes were more likely than others. This assumption, taken at face value, would seem to suggest that any given thinker has certain probabilities of having arrived at certain doxastic states toward a given proposition.

Taking all this on board allows us a natural way of thinking about what it means to be similarly reliable, independent, and an indicator of truth. Two thinkers are similarly reliable with respect to a given proposition just in case they have similar probabilities of getting it right (i.e. of forming true beliefs about the proposition). Two thinkers are independent with respect to P just in case the probability that the first thinker gets it right is equal to the probability that the first thinker gets it right *given* that the second thinker did.¹³ And finally, a thinker is an indicator of truth just in case the probability that she gets it right is greater than the probability that she gets it wrong.

We want to show that in the kinds of cases that interest Conciliationists, thinkers can be viewed as similarly reliable, and (more or less) independent indicators of truth. Start with similarity. These cases often involve disagreement between epistemic peers.

¹³ See Estlund (1994) p. 131 for a similar characterization. Also see Lackey (2013) for evidence that independence in belief-forming contexts is much more difficult to characterize precisely than it might first appear.

For our purposes, two thinkers are peers with respect to some proposition just in case they are equally reliable with respect to that proposition. That is, two thinkers are peers just in case they are similarly likely to believe the truth. Thus, it is natural to think that peers with respect to a given proposition also have similar probability distributions over the possible doxastic states they could take toward that proposition.¹⁴

Let's discuss independence next. It might be objected that the analogy breaks down here. While coins are paradigmatically independent of each other, it is doubtful that thinkers are. In many cases, thinkers tend to share the opinions of those around them more than one would otherwise expect. And occasionally, a contrarian thinker might be more likely to dispute a position if it is widely held than otherwise. If this is right, then it is hard to see how the analogy to coins goes through. In response, we want to make two points. First, in the cases of disagreement that Conciliationists discuss, the thinkers are not typically imagined to have based their judgments on each other's opinions. In light of this, it is harder to see how independence fails in these cases. Second, and more importantly, even if thinkers are not truly independent in these cases, we should note that independence comes in degrees. While pure statistical independence may be rare in real life thinkers, so is perfect statistical dependence.

¹⁴ Here we follow Elga (2007), White (2009), Enoch (2010), Kelly (2010), and Christensen (2014) in defining peerhood in terms of reliability. Other discussions in the literature, including Feldman (2007), Kelly (2005), Christensen (2007), and Cohen (2013), define peerhood in terms of *rationality*. Ultimately, either account of peerhood would suffice for our purposes. After all, rationality is presumably relevant to peerhood because it tends to be truth-conducive. Thus, it seems that equally rational thinkers would tend to be equally reliable, and therefore, have similar probability distributions.

Thinkers are not typically mindless sycophants or cartoonish contrarians. And it can make good sense to take a sample of partially independent opinions as a way to investigate truth. Imagine, for example, that the coins in the warehouse are not completely independent, but have a slight bias toward the side that the previous coin had landed. This failure of independence would not necessarily jeopardize one's ability to investigate whether P using these coins.¹⁵ Similarly, it seems that moderate departures from independence in real life cases of disagreement would not jeopardize one's ability to use peer opinion as a guide to whether P via sampling.

Of course, the partially independent coins are only useful if they are reliable indicators of truth. To defend the analogy, it must also make sense to believe that thinkers in ordinary cases of disagreement are indicators of truth. This is a minimum requirement for the opinions of those peers to be any evidence as to whether P. In general, it does make sense for one to believe this about one's peers. To see this, suppose that a thinker rationally believes P. To avoid a strange epistemic predicament, she must regard her epistemic peers as reliable indicators of the truth of P. Recall that if two thinkers are epistemic peers with respect to some proposition, then they have similar probability distributions over the possible doxastic states they can take toward

¹⁵ Suppose, for example, that the first coin has probability .75 of landing truly. Following a coin that lands truly, each coin has probability .9 of landing truly, and following a coin that lands falsely, each coin has probability .6 of landing truly. Suppose you take a sample of n flips and decide to believe the proposition that turns up most often. With $n=1$, your belief has a 0.75 probability of being true. With $n=3$, your belief has a 0.855 probability of being true. And with $n=5$, your belief has about a 0.910 probability of being true.

it. Thus, if a thinker were to believe that her peers were not indicators of truth – that is, that they were no more likely to be right than wrong – then it would be difficult for her to maintain that she is an indicator of truth herself. And there is something troubling about doubting one's own reliability in this way. At the very least, if a thinker cannot believe that her own belief is more likely true than false, then she suddenly has no epistemic reason to believe her own belief. Insofar as one is committed to P, there is pressure on one to regard one's peers (and oneself) as reliable indicators of the truth of P.

We can now see how the epistemic significance of peer opinion is similar to the epistemic significance of coin flips from the warehouse example. In the warehouse example, constructing samples of coin flips served as a means of investigating the truth of P. And larger samples were better than smaller samples. A parallel story can be told about peer opinion. Consider a math problem, with possible answers P and \sim P, about which you know that you and your friend are reliable indicators of truth. You think about the problem and come to believe P, while your friend does the same and concludes \sim P. Why should your friend's opinion give you reason to doubt P? Here's one way to reason. For independent, better-than-chance indicators of whether P, a sample of two is a better guide to P's truth than a sample of one. By previous argument, this case meets these conditions. This makes sense, as it's less likely for both thinkers to go wrong independently than for just one of you to do so. So you should certainly reduce

confidence in P – perhaps suspending judgment about P altogether. In general, it seems that a Conciliationist can think about the epistemic significance of peer disagreement in much the same way that one might think about the epistemic significance of coin flips in the warehouse example.

One might think that this is a strange way to think about Conciliationism. Why all this talk of probabilities and sampling? Actual people certainly don't think in these terms, and it doesn't seem like they are doing anything wrong in neglecting to do so. And indeed, powerful arguments for Conciliationism can be made without thinking about anything in statistical terms. In reply, we should note that we certainly don't think of our view as the uniquely correct way to think about disagreement. However, we do believe that there is nothing wrong with thinking about peer disagreement in this way – it does not lead to any mistakes. Thinking about Conciliationism in terms of sampling, however, does afford us a natural and elegant solution to the Modal Challenge.

4. Good and Bad Selection Procedures

We have seen that the epistemic significance of peer opinion can be understood in terms of sampling from probability distributions over what a peer might have believed. But not every opinion of a peer will give me evidence about which proposition I should believe. This is because not all samples provide evidence about the population or distribution from which they were drawn. How can we identify

problematic selection procedures and bad samples?

Let's return to the coin warehouse to see what a bad sample looks like.

Cherry Picking (Coin): You are standing outside the warehouse, with a coin in your pocket. You flip it once and obtain P. Your friend Nina enters the warehouse with the goal of finding coins that landed $\sim P$. After a few minutes, she emerges and hands you a list of one hundred warehouse coordinates and tells you that a coin that landed $\sim P$ at each location on the list.

Even if the locations on Nina's list do contain exclusively coins that landed $\sim P$, it should not make you think that $\sim P$ is true. Why not? Because you know that Nina 'cherry picked' certain outcomes with the goal of constructing a $\sim P$ -dominated list. Even if the vast majority of the coins in the warehouse landed P, Nina would be able to do this. And if Nina's list were constructed in this way, then we would have no reason to think that the sample she has furnished us with accurately represents the distribution of coin results inside the warehouse. In other words, we have reason to think that Nina has given us a biased sample. So we are entitled to dismiss the results of actual coin flips that disagree with ours because of how (we suspect) they were selected.

A natural thing to think is that Nina's sample was biased because its members were not chosen at random. Perhaps randomness is the selection constraint we seek? If it were, then we would have a problem, since ordinary cases of disagreement are not, generally, random samples. Randomness, however, is not necessary for a good sample. A convenience sample involves selecting whichever members of the population happen to be nearby. And convenience samples can sometimes be good ones. Suppose we went

into the coin warehouse and examined the one hundred coins closest to the entrance. So long as we had no reason to suspect that the coins had been sorted in advance, we would have every reason to trust the data we collected this way.¹⁶ Similarly, when I collect evidence about some proposition by surveying the opinions of nearby peers, this can have epistemic significance as a sort of convenience sample.¹⁷ So it seems that randomness cannot be the correct selection constraint.

Instead, the difference between Nina's sample and a sample taken from the one hundred coins nearest to the entrance is that only the latter should be taken to accurately reflect the coins' shared bias. Thinking about how these two samples were generated can illustrate this. Presumably, Nina constructed her list by admitting only those coins that turned up $\sim P$. Since you know this, you should think that the probability that a coin landed P , given that it was admitted to Nina's list, is effectively zero. This is much lower than the probability of any given coin's landing P . So learning that a given coin belongs to Nina's sample tells you something about how it fell. This is what is wrong with Nina's sample. And notice that this is not so with the convenience sample. Learning that a given coin belongs to this sample (i.e. learning that it was flipped near the door) tells you nothing about how it fell – that is, your confidence that it landed P , say, should be unchanged. This is why the convenience sample is epistemically significant.

¹⁶ In the table at the end of this section, we refer to this case as **Nearby Coins**.

¹⁷ In the table at the end of this section, we refer to this case as **Nearby Thinkers**.

This suggests a constraint on selection procedures for the coins. Such a procedure is good to the extent that you should believe that the probability that an arbitrary coin lands P (or \sim P) is approximately the same as the probability that an arbitrary outcome is P (or \sim P) given that it has been selected by the procedure.¹⁸ We can construct a similar constraint for selection procedures when we are sampling peer opinion rather than coins. For all members of a certain class of peers concerning some proposition P, there is some probability that a given thinker will come to believe P (or \sim P). A procedure is good to the extent that you have reason to believe that the probability that an arbitrary peer will believe P (or \sim P) is approximately the same as the probability that an arbitrary outcome is P (or \sim P) given that it was selected by the procedure.¹⁹

With this understanding of good selection procedures in place, we can return to cases of merely possible disagreement in order to apply the constraint. Let's start with a case of merely possible disagreement that does not seem to have epistemic significance.

Bare Possibility: You believe P. You realize that a peer could have believed \sim P.

You are trying to investigate P by constructing a sample of other opinions about P.

Setting aside your own opinion, you have a sample of one outcome: \sim P. In order for the

¹⁸ We can think of this as a conditional probability - the probability that a given outcome is P conditional on its having been selected.

¹⁹ The selection constraint that we posit here is only a necessary condition on a sample's having epistemic significance *as peer disagreement*. Obviously, a sample drawn from super-coins that were even more biased toward the truth would have epistemic significance as well - this is not relevant here. Additionally, at least one other constraint, that the members of a sample are statistically independent, is necessary for the sample's having epistemic significance in proportion to its size. These two conditions - the selection constraint and statistical independence - may well be jointly sufficient for samples of peer opinion to be epistemically significant, but our main point does not depend on this assumption.

sample to have epistemic significance, the procedure used must satisfy the constraint discussed above. Notice the procedure being used. You examine your own opinion, and then immediately consider the opposite opinion. This procedure effectively guarantees that you will arrive at a dissenting opinion rather than an agreeing one. The dissenting opinion, believing $\sim P$, was cherry picked from the set of attitudes an epistemic peer could have had. Simply by knowing that an opinion was selected by this procedure, we would know exactly what the opinion was. The probability that the opinion turned out to be $\sim P$ given that it was selected by this procedure is clearly higher than the probability that an arbitrary peer would believe $\sim P$. So this not a good selection procedure. Accordingly, you are justified in remaining steadfast in the face of the merely possible disagreement you encounter in this case.

Now let's think about the cases of merely possible disagreement that do seem to be epistemically significant – Coin Tyrant and Evil Tyrant. Recall:

Coin Tyrant: You flip the biased coin in your hand once and it lands P . You find other coins and you flip them. If Coin Tyrant realizes that a coin is going to land $\sim P$, he vaporizes it while it is in the air. This happens many times. You never observe any coin actually land $\sim P$.

Suppose you were to use your data to investigate whether P . Intuitively, the fact that many coins would have landed $\sim P$ seems relevant to whether you should believe P . The selection constraint on coins can be used to explain this. Suppose you opted to include only *actual* flips in your sample. Because of the tyrant's behavior, the probability that an arbitrary outcome is P , given that it was selected (i.e. that it was the result of a coin not

destroyed by the tyrant), is 1. Compare this value to the probability that an arbitrary coin lands P (i.e. the coin's bias with respect to P). In a certain strange sense, the probability that any coin lands P just is 1, since the tyrant consistently ensures that no coin ever lands otherwise. But this is not the relevant probability – it is not the bias. No matter the bias of the coin, the tyrant could ensure that no coin ever lands $\sim P$, or vice versa. Intuitively, what matters is not how a coin behaves in the presence of a P-obsessed tyrant, but rather, how it behaves under ordinary circumstances. And notice, if the coins truly had a bias of 1 toward P, then no vaporizations would have occurred under the tyrant's watch. So the probability of an arbitrary coin's landing P is far from 1 (as there were many such vaporizations), and for this reason, the selection constraint is violated. Looking only at actual outcomes won't suffice. On the other hand, if you were to expand the sample to include the would-have-been flips, then the probability that an arbitrary outcome is P given that it was selected by the procedure approaches the true bias of the coins.

Now recall the Evil Tyrant:

Evil Tyrant: A powerful, prophetic tyrant, decides that no one should ever believe the proposition P, so he orders the execution of all would-be P-believers (all of the people who would come to believe P, if they were permitted to live). 1,000 people are executed. Some time later, there is something of a consensus about whether P. The 1,000 remaining people all believe $\sim P$. There *could have been* disagreement about P if the tyrant had acted differently, but in fact there is only agreement.

Suppose you, a survivor, were to investigate whether P by surveying peer opinion.

Intuitively, the fact that many of your peers would have believed P, had they been

permitted to live, seems relevant to whether you should believe P. The selection constraint on thinkers can be used to explain this. Suppose you opted to include only actual opinions in your sample. Because of the tyrant's behavior, the probability that an arbitrary outcome is P given that it was selected by the procedure (i.e. given that it is the opinion of someone who was permitted to live) is 0. Compare this value to the probability that an arbitrary thinker believes P (i.e. the thinker's bias with respect to P). Just as with the coins, the fact that so many of your peers would have believed P had they not been killed makes it clear that your shared bias with respect to P is significantly greater than 0. So the constraint is violated, which implies that looking only at actual opinions does not suffice. On the other hand, if you were to expand the sample to include the would-have-been beliefs, then the probability that an arbitrary outcome is P given that it was selected by the procedure approaches the true bias of you and your peers.

Briefly, we would like to illustrate how this constraint applies to cases of actual disagreement. Interestingly, some cases of actual peer disagreement will turn out to be epistemically insignificant:

Cherry Picking: You take a math test with a million of your peers. You know that you all have a 90% success rate on problems like those on this test. After taking the test, you ask the teacher, who is reviewing the results, whether anyone's answer on question 17 disagrees with yours. She says "yes," and then produces a list of 100 peers who arrived at the other answer.

Intuitively, you should not be worried in this case, and the selection constraint can

accommodate this. Suppose that your answer was P and the only other possible answer was $\sim P$. Even if all of the test-takers are your peers, and even if one of them (or even one hundred of them) genuinely arrived at $\sim P$, you should not be worried for the same reason that you weren't worried about Nina's lengthy list in Cherry Picking (Coin). The selection procedure was biased. In this case, the probability that an outcome would be $\sim P$, given that it was selected by this procedure, is 1. This is, of course, much higher than the probability that an arbitrary peer would believe $\sim P$.

On the other hand, if you had simply asked the person sitting next to you about her answer on problem 17, and she had disagreed, you would have reason to be worried. This is a case of ordinary disagreement, and it does meet the constraint: The probability that an arbitrary outcome would be $\sim P$, given that it was selected by this procedure (i.e. given that the person seated next to you arrived at $\sim P$) is most certainly not 1. Indeed, this probability is equal to the desired probability – the probability that an arbitrary member of this group would arrive at $\sim P$. So conciliation is required here.

So the constraint applies to cases of actual disagreements as well as cases of merely possible disagreement. The following tables summarize the verdicts we have given in the cases we have discussed. Notice that the cases in which we should intuitively conciliate are precisely those that meet the constraint, independent of whether the disagreement is actual or merely possible.

BIASED COINS

	Good Sample	Bad Sample
Possible	Coin Tyrant	Bare Possibility (Coin)
Actual	Nearby Coins	Cherry Picking (Coin)

DISAGREEMENT

	Good Sample	Bad Sample
Possible	Evil Tyrant	Bare Possibility
Actual	Nearby Thinkers	Cherry Picking

6. Conclusion

Conciliationism threatened to collapse into an unacceptable skepticism because of the epistemic significance of merely possible disagreement. We have tried to show how Conciliationism can be maintained in the face of this threat. The explanation we have given involves understanding the epistemic significance of all peer opinion in terms of sampling. We take peer opinions that we encounter to be indicators of what we should think — provided that the opinions have been properly selected from the space of possible opinions we might have encountered. With this constraint established, we were able to distinguish Evil Tyrant from intuitively less threatening cases of disagreement. The less threatening cases all rely on an illicit sort of ‘cherry picking,’ which forced us to consider possibilities which we had no reason to consider very likely. While merely possible disagreement can have epistemic significance in special cases, it is important to note that such cases are rare in real life. And so, in the actual world, skepticism is avoided.

References:

Carey, Brandon (2011): “Possible Disagreements and Defeat,” *Philosophical Studies* 155: pp. 371-381.

- Carey, Brandon and Jonathan Matheson (2013): "How Skeptical is the Equal Weight View?" in *Disagreement and Skepticism* (ed. Diego Machuca), Routledge: pp. 131-149.
- Christensen, David (2007): "Epistemology of Disagreement: The Good News" *Philosophical Review* 116(2): pp. 187-217.
- 2011: "Disagreement, Question-Begging and Epistemic Self-Criticism", *Philosophers' Imprint* 11 (6) pp. 1-22.
- Christensen, David and Lackey, Jennifer (2013) *Disagreement: New Essays*. Oxford: Oxford University Press.
- Cohen, Stewart (2013), "A Defense of the (Almost) Equal Weight View," in D. Christensen and J. Lackey, eds., *The Epistemology of Disagreement: New Essays* (Oxford: Oxford University Press).
- Elga, Adam (2007): "Reflection and Disagreement," *Nous* 41: pp. 478–502.
- Enoch, David (2010): "Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement," *Mind* 119: pp. 953–997.
- Estlund, David (1994): "Opinion Leaders, Independence, and Condorcet's Jury Theorem," *Theory and Decision* 36 (2): pp. 131-162.
- Feldman, Richard (2007): "Reasonable Religious Disagreement," in L.M. Antony, ed., *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, (Oxford: Oxford University Press) pp. 194–214.
- Feldman, Richard and Warfield, Ted A. (2010) *Disagreement*, Oxford University Press.
- Frances, Bryan (2010): "The Reflective Epistemic Renegade," *Philosophy and Phenomenological Research* 81: pp. 419-463.

- Kelly, Thomas (2005): "The Epistemic Significance of Disagreement," *Oxford Studies in Epistemology* 1: pp. 167-196.
- (2010), "Peer Disagreement and Higher-Order Evidence," in Feldman and Warfield eds., *Disagreement* (New York: Oxford University Press).
- Kornblith, Hilary (2010): "Belief in the Face of Controversy" in *Disagreement* (eds. Richard Feldman and Ted Warfield), Oxford University Press, pp. 29-52.
- Lackey, Jennifer (2013): "Disagreement and Belief Dependence: Why Numbers Matter" in *The Epistemology of Disagreement: New Essays* (eds. David Christensen and Jennifer Lackey), Oxford University Press, pp. 243-267.
- Lasonen-Aarnio, Maria (forthcoming): "'I'm Onto Something!' Learning about the world by learning what I think about it" *Analytic Philosophy*.
- Lewis, David (1980): "A Subjective Guide to Objective Chance," *Studies in Inductive Logic and Probability* (ed. Richard C. Jeffrey), University of California Press, pp. 83-132.
- Sliwa, Paulina and Horowitz, Sophie (forthcoming): "Respecting *All* the Evidence" *Philosophical Studies*.
- Weatherson, Brian (ms.): "Do Judgments Screen Evidence?" available at <http://brian.weatherson.org/JSE.pdf>
- White, Roger (2009): "On Treating Oneself and Others as Thermometers," *Episteme* 6: pp. 233-250.