



Peer Disagreement and the Independence Principle

Epistemik Akranlar Arası Anlaşmazlık ve Bağımsızlık Prensibi

FAİK KURTULMUŞ

Sabancı University

Received: 15.06.2020 | Accepted: 27.06.2021

Abstract: This paper argues that the proper epistemic response to peer disagreement should depend, in part, on our grounds for judging the other person to be an epistemic peer. In particular, it argues that we can violate David Christensen's Independence principle if (a) our judgment that another person is an epistemic peer can only be established by track record judgments from the ground level, or (b) we put less trust in the reasons we relied on in judging the other person to be an epistemic peer than the reasons we relied on in arriving at the judgment that is the subject of our disagreement.

Keywords: Peer disagreement, independence principle, conciliationism, epistemic self-trust.



Introduction

When I believe that P and I find out that someone I consider an epistemic peer believes that not- P , or when I have high credence in P and an epistemic peer has low credence in P , what is the proper epistemic response to this fact? According to *conciliatory views*, once I find out that an epistemic peer has low credence in P my credence in P should be lowered. Or, if we adopt an all-or-nothing model of belief, when I find out that an epistemic peer believes that not- P and I believe that P , I should suspend judgment on P . There are cases where the conciliatory response is intuitively compelling. Here is one such case presented by Kelly:

You and I, two equally attentive and well-sighted individuals, stand side-by-side at the finish line of a horse race. The race is extremely close. At time t_0 , just as the first horses cross the finish line, it looks to me as though Horse A has won the race in virtue of finishing slightly ahead of Horse B; on the other hand, it looks to you as though Horse B has won in virtue of finishing slightly ahead of Horse A. At time t_1 , an instant later, we discover that we disagree about which horse has won the race. (Kelly, 2011, p. 184)

In this case, it is clear that my credence in Horse A winning should be significantly lowered. The conciliatory response is further strengthened by drawing analogies with the following kind of cases.

You and I are each attempting to determine the current temperature by consulting our own personal thermometers. In the past, the two thermometers have been equally reliable. At time t_0 , I consult my thermometer, find that it reads sixty-eight degrees, and so immediately take up the corresponding belief. Meanwhile, you consult your thermometer, find that it reads seventy-two degrees, and so immediately take up that belief. At time t_1 , you and I compare notes and discover that our thermometers have disagreed. (Kelly, 2011, p. 185)

The two cases seem analogous. Just as it would be wrong to take the fact that one of the thermometers is mine as a reason to favor its reading, it would be wrong to stick to my belief when we disagree about the outcome of the horse race. The conciliatory response might seem commonsensical. Yet, it would have radical implications. Think, for instance, of the big questions in philosophy. Suppose you think that humans have free will. Nevertheless, a lot of thoughtful and intelligent people think and



have thought that humans lack free will. The conciliatory view would require you to lower your confidence in your belief that we have free will even though you find the arguments for it persuasive. Given the widespread existence of peer disagreement about many of the central questions of philosophy, the conciliatory view would require us to lower our confidence in many of our philosophical beliefs. Thus, the proper response to peer disagreement is philosophically significant.

Despite the intuitive appeal of the conciliatory view, sometimes, when we find out that people we would consider to be epistemic peers do not share our responses this does not shake our confidence. We are inclined not to lose our confidence in our belief but demote others from the status of epistemic peers because of the disagreement between us. We are inclined to take the belief they have expressed to reveal a fact, not about the world but themselves. That is, we take the fact that they believe not- P when, by our lights, P is the case to establish that they are not as reliable as we are. In this paper, I shall argue that there are cases where this response is justified even in the absence of independent evidence favoring us. It is, ultimately, our grounds for having judged the other person to be an epistemic peer that determines whether this response is justified or not.

This response violates Christensen's Independence principle, which maintains that it would be wrong to demote someone's status as an epistemic peer in the absence of independent grounds. Thus, conciliatory views draw much of their support from this principle. Section 2 introduces the Independence principle. Section 3 distinguishes among different ways of determining others to be our epistemic peers. Section 4 argues that some ways of judging another person to be an epistemic peer justify violations of the Independence principle.

1. The Independence Principle

Let us define a person's epistemic peer as someone as reliable as that person in a certain domain.¹ That is, epistemic peers are equally likely to arrive at true beliefs in a given domain. More formally,

¹ There are several different definitions of epistemic peer in the literature. See for instance (Christensen, 2007, p. 188; Elga, 2007, p. 499; Kelly, 2005, pp. 174-5). In Section 3, I shall say more to defend my definition of epistemic peer.



Person A and Person B are epistemic peers on the question whether *P* or not-*P*, if and only if, when they are in circumstances *C*, (i) conditional on *P* being the case, the probability of A judging that *P* is equal to the probability of B judging that *P*, and (ii) conditional on not-*P* being the case, the probability of A judging not-*P* is equal to the probability of B judging not-*P*.

Circumstances include variables such as access to evidence and relevant arguments, and being in the same environmental conditions that have a bearing on the agents' reliability.

When I judge that *P* and someone I consider an epistemic peer judges not-*P*, there are three beliefs that are candidates for revision: my belief that *P* (or my credence *n* in *P*), my belief that the person I am disagreeing with is an epistemic peer, and my belief that we are in fact disagreeing.

Let us begin with the last candidate for revision. When I assert that *P*, I know that I am expressing what I take to be the truth. However, I do not know with the same certainty that you are also expressing what you take to be the truth. There is an important asymmetry between us regarding our knowledge of your intentions.² When you say that not-*P*, you may be lying or joking. So, my judgment that we are in fact disagreeing may be in error. In certain cases, believing that we are not actually disagreeing may be more sensible than believing that either my judgment that *P* or my judgment that you are an epistemic peer is in error.

My belief that the person with whom I am disagreeing is an epistemic peer can be revised in two ways. First, I might decide that even though you are an epistemic peer, on this occasion you are not. There may be an immediately available explanation of why, in this specific instance, you are more likely to be mistaken than I am. I might, for instance, think that on this occasion your cognitive faculties are not functioning properly because you have not had sleep for the past few days. Here we should note another asymmetry between our respective positions. I am in a position to monitor the functioning of my cognitive faculties more closely than I am in a position to monitor the functioning of your cognitive faculties (Lackey, 2010, p. 310).

Second, I might decide that you were not after all an epistemic peer,

² Christen also makes this point (2011, p. 9).



because you think not- P when P is the case. Those who are in favor of a conciliatory response to peer disagreement have been willing to accept the previous proposals but draw the line here. They hold that the previous responses are in agreement with the spirit of their view. What is crucial for the conciliatory response is that our response to peer disagreement respect what Christensen calls *Independence*:

In evaluating the epistemic credentials of another's expressed belief about P , in order to determine how (or whether) to modify my own belief about P , I should do so in a way that does not rely on the reasoning behind my initial belief about P . (Christensen, 2011, pp. 1-2)

The Independence principle does a good job of capturing what would be epistemically suspect in deciding that your thermometer was unreliable by looking at my thermometer in Kelly's example. It would be a question begging on my part to reason that your thermometer was unreliable because it disagreed with mine. Similarly, it captures what would be odd for me to decide that I need not worry about my disagreement with an epistemic peer, because P is true and, therefore, my epistemic peer is wrong.

Despite its intuitive plausibility and explanatory power, I shall argue that Independence is a principle we should *often, but not always*, respect. Whether we ought to respect it or not depends on the grounds of our belief that the person we are disagreeing with is an epistemic peer.

2. Judging another to be an Epistemic Peer

Other people's beliefs function like any other type of evidence in our deliberations about what to believe: we take them to be indicators of what is the case.³ You call me from Istanbul. You tell me that the weather there is cold. I come to believe that the weather in Istanbul is cold. I have built a gadget that relays me the readings of a thermometer in Istanbul. It reads 4 degrees Celsius. I come to believe that the weather in Istanbul is cold. Your belief functions in the same way as reading the thermometer. I take both your belief and the thermometer as indicators of the temperature in Istanbul. Insofar as I consider you a reliable indicator of the temperature, I will give weight to your beliefs about the temperature.

³ Kelly makes this point (2011, pp. 196-7).



How do I decide whether you are reliable or not? In many cases, we assume that people are reliable indicators of various states of the world they report on. Stopping to determine the reliability of others in every instance would have a paralyzing effect on our lives. In many domains, we treat others as reliable unless proven unreliable. Call this *presumed reliability*. When I ask you about the weather in Istanbul, I will assume that (a) you will answer me truthfully; and (b) your answer is likely to be correct. Under normal circumstances, I would not try to determine whether you are reliable. Presumed reliability is a very rough judgment. Moreover, we presume that people are reliable, not that they are equally reliable. For these reasons, it cannot be used to establish that two people are epistemic peers.

Not all judgments of reliability are presumed. There is also, what I will call *inferred reliability* where we do not presume that a person is reliable but determine that they are so. There are two main ways of determining someone's reliability. First, I can judge that they are reliable by examining their track records. Second, my background theories may suggest that they are reliable. Suppose I wonder whether you are reliable at doing sums in your head. I give you a set of numbers. You add them up in your head. I do the sums on paper and check my results twice. We repeat this procedure several times. Insofar as you come up with the results I come up with, I'll conclude that you are reliable. In this case, I have what I take to be a very reliable way of checking the results you have produced. Note that, in this instance, when judging whether you are reliable at doing mathematical operations in your head, I am not assuming that I am reliable at doing mathematical operations in my head. I am assuming that the method of doing the sums on paper and checking the results is a reliable method. I could examine whether I am reliable by using the same method and come to decide that we are equally reliable. In judging someone else's reliability in this instance, we took up an epistemically more secure position than them. Doing sums on paper and checking it twice is a more reliable way of doing sums than doing them in one's head. I shall refer to these kinds of judgments as *track record judgments from a more secure epistemic position*, because we establish others' and our reliability by an assessment of our track records in light of judgments that are more reliable.



Sometimes, however, we do not have the opportunity to judge someone's reliability by taking up a more secure epistemic position than them. In those cases, we have to decide whether they are reliable by comparing what they take to be the case with what we take to be the case without taking up a more secure epistemic position. I shall call these judgments *track record judgments from the ground level*.

Here is an example of a track record judgment from the ground level. Suppose I am an art historian. I hear in the news that an art historian claims that a drawing attributed to Michelangelo is a forgery while another art historian denies this. I am not familiar with either art historian, and I lack knowledge of their credentials. I cannot study the drawing myself. However, as luck would have it, there are several drawings attributed to Michelangelo where I live, and both art historians have studied them. There's a record of which ones they thought were genuine and which ones they thought were forgeries. In addition to this, I know that none of these studies involved any special instruments. They were just done by close observation. Lacking instruments that would enable me to make track record judgments from a more secure epistemic position, I can look at the drawings and compare the judgments they have made with mine to determine which art historian to trust. If one of them seems consistently wrong, whereas the other seems consistently right, then I should trust the latter. Furthermore, if we are in substantial agreement about the drawings I should judge her to be an epistemic peer. In this example, my judgment of reliability is not made from an epistemically more secure position. I am in the same position as the two art historians. None of us had access to instruments that would enhance our reliability. I am judging their reliability by examining how well their judgments match up with what I take to be the case.⁴

Sometimes I may have good reason to assume that someone is reliable even though I do not have the opportunity to examine their track record. In such cases, we can make *theory-based judgments of epistemic peerhood*. Background theories I hold may require me to judge that they are as reliable as or more reliable than I am in a given domain. For instance, given facts

⁴ Track record judgments from the ground level are similar to what Kitcher, in a different context, calls direct calibration (1993, pp. 316-8).



about the physical make up of your eyes, how the human eye works, and the outside conditions necessary for good vision, I can conclude that your vision reports ought to be reliable without any reference to these reports. Coupled with knowledge of similar facts about me, I can find out that we should be equally reliable.

I should at this point remark on why I have defined an epistemic peer as I did. As you will recall, I defined an epistemic peer as someone who I consider to be as reliable as I am in a certain domain. There are other definitions in the literature. For instance, according to Kelly

Two individuals are epistemic peers with respect to some question if and only if they satisfy the following two conditions: (i) they are equals with respect to their familiarity with the evidence and arguments which bear on that question, and (ii) they are equals with respect to general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias. (Kelly, 2005, pp. 174-5)

Those two individuals satisfy the conditions set out by Kelly, is a reason to expect that they will be equally reliable. Nevertheless, what matters is that they are equally reliable in the way I defined above. If the conditions laid out by Kelly did not contribute to one's likelihood of having true beliefs, why would we care about them? Alternatively, if someone can be reliable even though she does not satisfy Kelly's conditions, why should we not treat them as epistemic peer? Suppose there is a person who you, in general, find to be thoughtless, and careless in her reasoning. However, when, for instance, doing sums in her head she gets the right answer as often as you do. That person should count as an epistemic peer in that domain, and her answers should have a role in your reasoning even though she does not satisfy Kelly's conditions for epistemic peerhood. She is a good indicator of what is the case in that domain. Kelly's, and other similar definitions, are better interpreted as background theories which would give us a reason to expect two people to be epistemic peers rather than as a definition of epistemic peers.

3. Judgment of Reliability and Response to Disagreement

After this survey of the ways in which we come to judge others to be our epistemic peers, we can re-examine the Independence principle. In



cases where I have inferred your reliability from your track record by comparing what you take to be the case with what I take to be the case without occupying a more secure epistemic position, I should be entitled to violate Independence. For suppose that it would be wrong in those cases where I have made track record judgments from the ground level to violate Independence. Then, the order in which you and I expressed our judgments would be significant. But this is unacceptable.

To see this, suppose that we both express our judgment on a set of questions, $q_1; q_2; \dots; q_i; \dots; q_n$. Let's say that our answers to $q_1; \dots; q_i$ are in agreement. By looking at our track records I should consider you an epistemic peer. Now suppose that our answers to questions $q_i; \dots; q_n$ are not in agreement. The Independence principle prohibits me from demoting you in the absence of independent grounds, which, in this instance, I lack. Thus, if I am not allowed to violate Independence after I have decided that you are an epistemic peer, then I am required to change my mind about my answers to $q_i; \dots; q_n$ since an epistemic peer disagreed with my answers. If we began the sequence of questions from q_n instead of q_1 , I would not have identified you as an epistemic peer, because we disagreed on many questions. Accordingly, I would not be required to change my mind about my answers to $q_i; \dots; q_n$. However, the order in which we begin the sequence of questions should not make a difference to what it is rational for me to believe. It should only depend on the total evidence I have.

Let us go back to our art historians' example to illustrate this problem in a case. Suppose that I judged one of the art historians to be as reliable as I am. Then, one day I notice that there was another set of Michelangelo drawings I had access to, and this art historian has made judgments on their authenticity. I find that all of our judgments are in disagreement. If there is no independent reason for me to believe that his judgments on these occasions were less reliable, I am required to change my mind about these drawings in order not to violate Independence. However, if, in the past, I had not decided that this art historian was as reliable as I am, I could use my opinions about this new set of drawings to judge that this art historian was not reliable.

It should be clear from these examples that, insofar as we are justified in relying on track record judgments from the ground level to determine



that another person is an epistemic peer, and this is the only basis of our judgment that we are epistemic peers, we should be entitled to violate Independence. Otherwise, what is rational for us to believe in light of the same set of evidence would be sensitive to the order in which we compare our judgments with another person.

One possible response to this argument is to hold that it goes wrong in supposing that we can make a judgment regarding the other person's reliability before all the evidence is in. So, we should make a judgment only when we've compared judgments on $q_1...q_n$. The trouble with this response is that it will rob the Independence principle and the conciliatory response to peer disagreement of all their teeth. All instances of agreement and disagreement are relevant as evidence. If we are required to wait for all the evidence to come in before we judge that another person is an epistemic peer, we will have few, if any, opportunities to put the Independence principle into practice.

Another possible response would be to challenge the propriety of relying on our judgments from the ground level to determine whether someone is reliable in the way that we did in the art historian's example. The natural response to this challenge is that we have no choice but to rely on our judgments when a more secure epistemic position or a background theory is unavailable.⁵ Who else's judgments should we rely on? And if we are to rely on someone else's judgments, how are we to decide whose judgments we should rely on? We are inevitably led back to our own judgments as the final arbiter.

This response needs to be qualified. In certain cases, there have to be intermediary steps before the appeal to judgments from the ground level. Otherwise, we would be epistemically irresponsible. Consider the test I suggested for determining whether someone is reliable at doing mathematical operations in their head. I suggested that we could do the operations on paper and check our results twice. Given that this method was available, it would be epistemically irresponsible of me to test your reliability by comparing the results you obtained when you did the mathematical operations in your head with the results I got when I did them in my head. Similarly,

⁵ David Enoch also makes this point (2010).



if there was a well-supported theory that spelled out the conditions for when a person is going to be reliable in doing mathematical operations, relying on track record judgments from the ground level would be epistemically irresponsible. When there is the possibility of independent checks on our results and judgments we should rely on them. But ultimately some sources of knowledge admit of no independent checks. In those cases, we are entitled to rely on our judgments, which we make without taking up an epistemically more secure position, to determine that another person is an epistemic peer. After all some level of epistemic self-trust, in the sense of reliance on one's cognitive capacities without prior justification of their reliability, is inescapable (Alston 1993; Foley 2001; Zagzebski 2012).

Conclusion

I conclude that we are entitled to rely on track record judgments from the ground level, and when our judgment that the other person is an epistemic peer is based on such a method, we should be allowed to violate Independence. If judgments from the ground level can be used to determine that another person is our epistemic peer, they can also be used to undermine this initial judgment.

That we be should be allowed to violate Independence in these cases does not entail that it is always rational to demote others when we disagree with them. We leave some margin for error and invest different degrees of confidence in our judgments even when our judgments admit of no independent checks. If I identify another person as an epistemic peer on the basis of track record judgments from the ground level, and we disagree on something which I am not very confident of, I should be reluctant to demote them from the status of epistemic peer because of this disagreement. Similarly, if our past agreements were numerous and significant, I should be reluctant to demote you from the status of epistemic peer. Since I am not infallible, there will be cases where I will disagree with an epistemic peer and be in the wrong. The rejection of the Independence principle does not require us to deny these points. What the rejection of the Independence principle entails is that the disagreement itself can act as evidence in two ways, even when we lack independent grounds: (a) it may give us reason to think that we may be wrong; or (b) it can give us reason to



think that the other person is not an epistemic peer despite our initial assessment to the contrary.

Should we respect Independence when our reliability judgment is based on background theories? The key point to keep in mind, throughout, is that we have no choice but to rely on our judgments and that our judgment that another person is an epistemic peer is fallible like our other judgments. The intuitive appeal of the Independence principle and conciliatory views lies in the very good question ‘What’s so special about you?’ No doubt, this question has force. However, from the first person-perspective there is another equally forceful question that needs to be asked: ‘What’s so special about my judgment about the reliability of my interlocutor?’ The answer to our question whether Independence needs to be respected in the case of reliability judgments based on background theories, therefore, depends on our confidence in the background theory and how strongly our background theory supports the judgment that the other person is an epistemic peer. When we are confident of our background theory, and of the judgment that it gives us reason to identify the other person as an epistemic peer, we have good reason to respect the Independence principle. This point is well-illustrated by the Horse Race case. In that case, we have good reason to think that our judgment may be mistaken, and our background theory gives us good reason to believe that the other person is an epistemic peer. When, however, we are not very confident of our background theory, or we are not confident that our interlocutor satisfies the conditions that would, according to our background theory, make them our epistemic peer, we may violate the Independence principle.

In the case of theory-based judgments of epistemic peerhood, the Independence principle draws support from the apparent epistemic symmetry between us and our epistemic peer. The thought is something like the following. My epistemic peer and I satisfy certain conditions, *C*. *C* is the main determinant of one’s likelihood of arriving at true beliefs. Therefore, I should expect my epistemic peer to be as likely as I am to have true beliefs. The Independence principle, quite plausibly, suggests that when we are confident of the fact that we both satisfy *C*, and of the fact that *C* is the main determinant of one’s reliability, then we should not rely on the belief under contention to demote the person we disagree with from the



status of epistemic peer. However, when we are more confident of the judgment under contention than these premises, violating Independence is more reasonable, because our belief in the existence of epistemic symmetry between us is weaker.

Should we be allowed to violate Independence when our knowledge of another's status as an epistemic peer is based on our evaluation of their track record from a more secure epistemic position? In such cases, the previous argument, which I gave to show that we should be allowed to violate Independence when our track record judgments are from the ground level, does not apply. Moreover, in the case of track record judgments from a more secure epistemic position, our judgment that *P* would be less reliable than our judgment that we are epistemic peers. Accordingly, we should respect Independence in these cases.

I would like to address two final worries. The position I have defended that allows me to demote others from the status of epistemic peerhood because they disagree with me may seem epistemically arrogant. The following thought, however, should mitigate the appearance of epistemic arrogance. When I decide that you were not, after all, an epistemic peer my reasoning is not that I am a better epistemic agent, so I am always entitled to discount your opinion. Rather, I think that my judgment that we are epistemic peers is more likely to be wrong than my judgment that *P*. Note also that I do not need to assume that my judgment that *P* cannot be mistaken. It is just that my belief that you are an epistemic peer is more likely to be mistaken. Even though this line of reasoning displays epistemic arrogance in one respect, it shows epistemic modesty in another.

When I demote another person from the status of epistemic peerhood, is it not possible that I am making a mistake and foregoing a chance to improve my beliefs? There is, of course, this possibility. But there is also the possibility of mistakenly treating another as an epistemic peer. The best that we can do is to be guided by what we take to be the best reasons we have. Judging that another person is not an epistemic peer is no less risky than judging that another person is an epistemic peer.



References

- Alston, W. P. (1993). *The Reliability of Sense Perception*. Ithaca: Cornell University Press.
- Christensen, D. (2007). Epistemology of Disagreement: The Good News. *The Philosophical Review*, 116 (2), 187-217.
- Elga, A. (2007). Reflection and Disagreement. *Noûs*, 41 (3), 478-502.
- Enoch, D. (2010). Not Just a Truthometer: Taking Oneself Seriously (But Not Too Seriously) in Cases of Peer Disagreement. *Mind*, 476, 953-997.
- Foley, R. (2001). *Intellectual Trust in Oneself and Others*. Cambridge: Cambridge University Press.
- Kelly, T. (2005). The Epistemic Significance of Disagreement. *Oxford Studies in Epistemology*, vol. 1. (Eds. T. S. Gendler & J. Hawthorne). Oxford: Clarendon Press, 167-196.
- Kelly, T. (2011). Peer Disagreement and Higher Order Evidence. *Social Epistemology: Essential Readings*. (Eds. A. I. Goldman & D. Whitcomb). Oxford: Oxford University Press, 183-217.
- Kitcher, P. (1993). *The Advancement of Science: Science without Legend, Objectivity without Illusions*. Oxford: Oxford University Press.
- Lackey, J. (2010). A Justificationist View of Disagreement's Epistemic Significance. *Social Epistemology*. (Eds. A. Haddock, A. Millar, & D. Pritchard). Oxford: Oxford University Press, 298-325.
- Zagzebski, L. T. (2012). *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief*. Oxford: Oxford University Press.

Öz: Bu makale, epistemik akranlarımızla anlaşmazlığa düştüğümüzde epistemik açıdan takınmamız gereken tavrın karşımızdakinin epistemik akranımız olduğunu düşünmek için sahip olduğumuz nedenlere bağlı olduğunu savunmaktadır. Literatürde David Christen'in savunduğu Bağımsızlık Prensibi'ni eleştiren makale, bu prensibe aykırı davranmanın hangi durumlarda epistemik açıdan uygun olduğunu incelemektedir.

Anahtar Kelimeler: Epistemik akranlar arası anlaşmazlık, bağımsızlık prensibi, uzlaşmacılık, epistemik öz-güven.

