


ARTICLE

Conservative Treatment of Evidence

Alireza Fatollahi 

Bilkent University, Ankara, Turkey

Email: alireza.fatollahi@gmail.com

(Received 19 February 2022; revised 14 June 2022; accepted 5 July 2022)

Abstract

This paper discusses two conservative ways of treating evidence. (I) *Closing inquiry* involves discounting evidence bearing on one's belief unless it is particularly strong evidence; (II) *biased assimilation* involves dedicating more investigative resources to scrutinizing disconfirming evidence (than confirming evidence), thereby increasing the chances of finding reasons to dismiss it. It is natural to worry that these practices lead to irrational biases in favor of one's existing beliefs, and that they make one's epistemic condition significantly path-sensitive by giving a bigger role to batches of evidence obtained earlier in the course of inquiry compared with those subsequently acquired. However, I argue that both practices are demanded by considerations of practical rationality. I also argue that, contrary to initial appearances, there is little reason to worry about the effects of these practices on the dynamics of one's beliefs.

Keywords: Biased assimilation; closing inquiry; conservatism; practical rationality; bias

1. Introduction

What ought one to do when one obtains purported evidence bearing on one's beliefs? Here is an initially plausible answer: one ought to investigate such evidence to determine whether it is indeed genuine evidence. If so, one ought to proportion one's beliefs to the totality of one's evidence, which includes the newly acquired piece. This is an attractive picture, but it is an idealization. We obtain too many pieces of purported evidence for us to thoroughly investigate all of them. Thus, we must inevitably prioritize. In this paper, I am interested in two conservative ways in which one can approach such prioritization of one's investigative resources in face of purported evidence against one's beliefs. The first conservative approach is what I call *closing inquiry*. It involves discounting purported evidence bearing on a belief as misleading or unreliable evidence without much or any critical scrutiny, unless the evidence is strong enough to constitute "positive" reason for re-opening one's investigation of that belief.^{1,2} By discounting

¹This is an aspect of what Gilbert Harman (1986: Ch. 5) has called "acceptance as full belief," which involves believing that p and closing one's investigation of p . According to Harman, only strong enough *counterevidence* can rationally require one to re-open one's investigation in a belief. However, I think strong confirming evidence can sometimes have the same effect. If p is an important belief, one might acquire strong enough evidence for it that would rationally warrant (or demand) re-opening one's investigation of p just to gain certainty.

²Closing inquiry must be distinguished from the kind of dogmatic practice discussed originally by Saul Kripke in connection with his "dogmatism paradox." (This appeared in 'On Two Paradoxes of Knowledge,'

(or dismissing) purported evidence, I mean only setting it aside without necessarily forgetting it.³ The second conservative approach involves dedicating more cognitive resources to critically examining counterevidence (compared with confirming evidence),⁴ and thereby increasing the chances of finding grounds for questioning its credibility. In social psychology literature, this is called *biased assimilation* of data.⁵

Both practices appear to be grounded in an irrational resilience against change in one's beliefs. Moreover, they appear to have at least two worrisome effects. First, they seem to work as biases in favor of one's beliefs. And second, they appear to render one's epistemic condition path-dependent, by placing more emphasis on the first batches of evidence one obtains compared with evidence subsequently acquired.⁶ This is perhaps best manifested in the well-documented causal linkage between biased assimilation and "belief polarization." This phenomenon occurs when people who begin with initially opposing views become even more divided upon exposure to identical bodies of evidence parts of which confirm one party to the dispute, while its other parts confirm the other party. By dedicating more cognitive resources to scrutinizing pieces of disconfirming evidence, one increases the chances of finding ways for explaining them away as misleading evidence. This in turn might lead to a credence-boost in one's initial belief after acquiring evidence of such mixed character. Thus, if two parties who disagree on an issue exhibit biased assimilation towards an identical body of mixed evidence, they might end up disagreeing more strongly than they initially did.⁷ In light of such worrisome consequences of biased assimilation, Kelly maintains that "those few of us who are aware of the phenomenon of belief polarization ... ought to be less confident of beliefs that are likely to have benefited from the underlying psychological mechanisms [biased assimilation]" (Kelly 2008: 629).

In this paper, I examine these worries and conclude that they are largely unwarranted. The first part of the paper examines the normative foundations of closing inquiry and biased assimilation: under what conditions, if any, is it rational to engage in them? The second part explores the effects they tend to have on the dynamics of one's beliefs: whether and to what extent they tend to render one's epistemic condition path-dependent.

In section 2, I will discuss the notion of consequentiality of purported evidence, which is defined as follows.

an unpublished lecture delivered to the Cambridge Moral Sciences Club.) A Kripkean dogmatist discards any evidence against a belief she already holds as misleading evidence *and* will not re-open her investigation in her belief regardless of its strength.

³In section 3.2, I will get back to the complications that forgetting evidence introduces to my discussion.

⁴I am assuming that such disparity in treating confirming and disconfirming evidence is only in terms of dedicating one's resources to their scrutiny. Presumably, even then one's epistemic standards (for example, whether and to what degree a piece of evidence confirms/disconfirms a given hypothesis) are the same for confirming and disconfirming evidence. Otherwise, I find such a practice epistemically irrational. (They might nonetheless be *practically* rational. See Shaffer (2019).)

⁵The term 'biased assimilation,' as it is understood in that literature, has other meanings too. For example, it could mean "a propensity to remember the strengths of confirming evidence but the weaknesses of disconfirming evidence." Here I use the term only to refer to the practice described above. Lord *et al.* (1979) is the classic study on the topic. See also Kelly (2008) for a philosophical examination.

⁶Kelly (2008: 628–9) has discussed this worry with respect to biased assimilation. I will make similar observations in connection with closing inquiry in section 3.

⁷The mechanism by which biased assimilation leads to belief polarization is explained in greater detail in section 3.

Consequentiality of a piece of purported evidence, *e*, for a rational agent is defined as the amount of change she must make in the totality of her epistemic attitudes in order rationally to accommodate *e*.

I begin by defending the idea that the consequentiality of a piece of purported evidence is an important factor in determining the appropriate amount of resources one ought rationally to dedicate to scrutinizing it. I then argue that this idea justifies closing inquiry and biased assimilation, each in its own appropriate conditions. In brief, the rationality of closing inquiry lies in the fact that when a belief is highly probable, evidence bearing on it is relatively inconsequential. And the rationality of biased assimilation rests on the fact that other things being equal, evidence against one's belief is more consequential than evidence for it.

In section 3, I discuss how biased assimilation and closing inquiry tend to influence the dynamics of one's beliefs. I argue that, initial appearances to the contrary, in the long run neither practice adversely affects one's epistemic condition – at least not to an excessive degree. Closing inquiry tends to have no substantial effect on one's epistemic conditions in the long run if it works in isolation. And interestingly, biased assimilation may work as a bias *against* one's beliefs in the long run. I end section 3 by observing that there are otherwise unproblematic epistemic traits, such as forgetting evidence, that when combined with biased assimilation or closing inquiry might lead to more worrisome belief-dynamics in the long run than these practices tend to do on their own.

A few important clarifications are in order at the outset. First, biased assimilation and closing inquiry are modes of *inquiry*. Although what I say in this essay has some bearing on the normativity of *beliefs* acquired through these practices, that issue is not my main focus and my discussion of it is limited to a few isolated remarks in section 3. Importantly, the practices that I talk about are different from other conservative approaches to counterevidence that involve violating norms of epistemic rationality.⁸ Indeed, as I shall argue, even impeccably epistemically rational agents might exhibit biased assimilation or closing inquiry in certain conditions. Second, closing inquiry is neither necessary nor sufficient for outright belief. Closing one's investigation of *p* is a mode of *treating* purported evidence bearing on *p*, while believing that *p* seems (to me at least) to be a matter of the epistemic standing of *p*.⁹ Third, the position I defend in this paper is compatible with both permissivism – the thesis that more than one doxastic attitude might be rational to hold given one's total evidence – and its denial.¹⁰ This is because biased assimilation and closing inquiry can only make a difference to what one's total evidence is. They are simply silent on how one ought to proportion one's beliefs to one's total evidence once that evidence is fixed.

2.1. Consequentiality-based inquiry

Any treatment of *genuine* evidence other than proportioning one's beliefs to it is irrational. However, not all purported evidence is genuine evidence.¹¹ Suppose you

⁸An important example of such practices is what Michael Shaffer has called “motivated pragmatically rational epistemic irrationality.” See Shaffer (2019) for an interesting discussion.

⁹For a few representative views on how outright belief is related to partial belief, see Foley (1993), Frankish (2009) and Leitgeb (2014). My argument doesn't presuppose any particular view in this literature.

¹⁰It has been suggested that Harman's version of conservatism (including his account of acceptance as full belief) leads to epistemic permissivism (White 2005: 445–6).

¹¹Although the distinction between purported and genuine evidence is crucial to my argument, I won't offer a formal account of this distinction. This is because any formal treatment of the distinction, as far as I

read in a relatively serious news outlet with strong political inclinations that ‘there’s been a peaceful protest in city X by people whose political views are similar to those of the news outlet’ (p). You have *purported* evidence that p . But you might be concerned that the news outlet has downplayed possible violent aspects of the protest. If this concern undermines your confidence in the report seriously enough, you might decide not to accept it as genuine evidence and discount the report or seek further evidence either directly about p (by, say, reading reports of the incident in other reputable news outlets with different or weaker political tendencies) and/or about the original piece of evidence (for example, by inquiring about the particular journalist whose report you’ve read, etc.).¹² Another example. The results of scientific studies are not always immediately treated as genuine evidence by scientists. Sometimes they obtain that status only if they can be reproduced or proven “compatible” with the general trends in larger bodies of data (i.e., after serious critical scrutiny).

In general, one has three options in face of purported evidence. One might simply accept it as genuine evidence without further scrutiny; one might scrutinize it before accepting or discounting it; or one might outright dismiss it. The decision how to treat purported evidence is significantly influenced by one’s level of confidence in it being genuine evidence (relative to the relevant standards in the theoretic task in question). And one’s confidence is determined by one’s higher-order evidence – evidence about the nature and bearing of the original piece of evidence on one’s beliefs, given one’s epistemic condition.¹³ In the rest of this paper, I will consider only those pieces of purported evidence that, given one’s (lack of) confidence in them being genuine evidence, one cannot accept as genuine evidence without further scrutiny. (I have nothing interesting to say about purported evidence one finds genuine other than that one ought to proportion one’s beliefs to it.) When it comes to this kind of purported evidence, one faces the choice of how much investigative resources, if any, to dedicate to it. Here again one’s confidence in the evidence being genuine is important. Other things being equal, if one is only slightly less confident of its being genuine than what is required for accepting it, one is more likely to investigate it (as opposed to discount it) than if one is seriously doubtful about it. However, this is not the only relevant factor. Given the fact that one’s resources are limited, to decide to scrutinize a given piece of purported evidence is to decide not to scrutinize another piece or, say, not to read a novel. Thus, the rationality of engaging in such a scrutiny is influenced by what one expects to gain from it. And here the consequentiality of evidence is key.

The crucial fact on which the relevance of consequentiality of evidence rests is that evidence has a dual significance for an agent with the epistemic goal of believing truths and not believing falsehoods. On the one hand, the agent is bound by the demands of epistemic rationality: she ought to proportion her beliefs to her evidence. On the other

can tell, will be based on, or at least closely linked to, a particular view of evidence. (And the existing views on what constitutes evidence are widely different, even by standards of philosophy.) Instead, I hope to make the distinction sufficiently clear by a few examples.

¹²One might point out that in this example you have genuine (not merely purported) evidence that ‘you’ve read a report that p ’. This might be true but it’s beside the point. This evidence bears on the rationality of your belief that p inasmuch as it gives you genuine evidence that p . Thus, the interesting question here is not whether I have genuine evidence that ‘I’ve read a report that p ’ but rather whether I have genuine evidence that p .

¹³There is an extensive literature on the role of higher-order evidence in the proper assessment of evidence. See for example, Feldman (2009), Christensen (2010), Kelly (2010) and Lasonen-Aarnio (2014). My argument doesn’t presuppose any particular view in that literature.

hand, evidence is instrumental towards her epistemic goal of finding truths:¹⁴ acquiring new bodies of evidence typically improves the agent's epistemic condition. If she has unlimited resources, she ought to try to acquire all the evidence she can (first-order and higher-order), and then proportion her beliefs to the totality of such evidence. However, if she has limited resources, she ought to allocate them to an examination of this rather than that piece of purported evidence (or rather than, say, play chess) based on how significant she expects the outcome of that scrutiny to be towards her epistemic goal and also based on the relative importance of the epistemic goal itself. Now since she cannot know the outcome of the investigation beforehand, she ought to decide on the basis of her *expectations* of its outcome. Such expectations are determined by a number of factors. However, a plausible estimate for the agent's expectations of an examination of a purported piece of evidence, *e*, is the expected difference (in her view) between scenario (a): she finds *e* genuine evidence and proportions her beliefs to it, and scenario (b): she finds *e* misleading evidence and discounts it. Here the expected difference is determined by (i) the agent's subjective *probability* that (a) occurs instead of (b) and (ii) how different the totality of her epistemic attitudes ought to be if (a) occurs instead of (b). But this expected difference between (a) and (b) is the expected consequentiality of *e* (see the definition of consequentiality on page 3). Therefore, we have the following key fact.

Key Fact: The amount of resources a rational agent ought to allocate to an examination of a piece of purported evidence is determined, among other things, by the expected consequentiality of that evidence for her.

I will shortly argue that Key Fact underlies the justification for conservative treatments of purported evidence under certain conditions. Before doing so, however, I must make a few observations about the notion of consequentiality. First, I have defined consequentiality in terms of 'epistemic attitudes,' instead of, say, credences or outright beliefs, because considerations of consequentiality are relevant to the rational allocation of investigative resources independently of how fine-grained one partitions one's doxastic attitudes.

Second, consequentiality is a feature of purported evidence, not evidence one already takes to be genuine. In particular, the consequentiality of a piece of evidence is not solely determined by how much change in one's doxastic attitudes rationally accommodating it would require *in the event that one finds it genuine evidence*. Another factor is whether it *is* genuine evidence. Thus, an important factor in determining the expected consequentiality of a piece of purported evidence is the subjective probability of its being genuine evidence. For example, if I tell you something absurd but amazing if true about the nature of reality, the difference between scenario (a) and scenario (b) for this piece of testimonial evidence is large; however, if you find it extremely unlikely that what I told you is true, the expected consequentiality of this evidence is low and you can rationally discount it without investigation.

Third, if E_1 and E_2 are two bodies of evidence that bear on H , then, other things being equal, E_1 is more consequential than E_2 if rationally accommodating E_1 requires

¹⁴Epistemologists disagree on whether epistemic rationality is instrumental rationality in the service of an epistemic goal (such as "now believing truths and now not believing falsehoods") or whether it is not a type of means-end rationality at all. See for example Foley (1987), Feldman (2000) and Kelly (2003). My claim here is compatible with both views on this issue.

a bigger change in one's epistemic attitude towards H . However, the effect of accommodating, say, E_1 will not be limited to one's attitude towards H alone. A change in one's attitude towards H will have a ripple effect on one's other beliefs. And the more central H is to one's "web of belief," the more severe such ripple effects will typically be, by requiring revisions in a larger number of beliefs and/or by involving more severe changes in other beliefs.¹⁵ In fact the consequentiality of a piece of evidence cannot be represented by a single number, because for two arbitrary bodies of evidence, E_1 and E_2 , it is *not* the case that either E_1 is more consequential than E_2 or E_2 is more consequential than E_1 or they are equally consequential. E_1 might require a larger change in one's attitude towards p than E_2 requires towards q , while a change in one's attitude towards q would have an effect on a larger number of other beliefs than a similar change in one's attitude towards p . Fortunately, such complexities of the notion of consequentiality will not be relevant to my discussion.

Fourth, consequentiality of evidence captures *only one* aspect of the all-things-considered importance of evidence. Other epistemic and practical factors are also relevant to the all-things-considered importance of a piece of evidence. Consequentiality is a function of what James Joyce has called "the balance of evidence," where the overall balance of a body of evidence bearing on a proposition is a matter of how decisively it tells in favor of that proposition. However, evidence can have other epistemically important features, e.g., in terms of its "weight" or "specificity." For example, suppose I believe a coin is fair and I obtain the following pieces of purported evidence. E_1 : the coin was tossed a million times and it landed heads exactly half of the time. E_2 : the coin was tossed 10 times and it landed heads 9 times. E_2 is more consequential because E_1 doesn't require any change in my belief that the coin is fair. However, E_1 is weightier. It results in a more "stable" or "resilient" credal state.¹⁶ I think in this particular example E_1 is all-things-considered more important. Thus, if I can scrutinize only one of E_1 or E_2 (and other things are equal), I must scrutinize E_1 . The all-things-considered importance of a piece of evidence is affected by non-epistemic facts of the situation too. Consider: I currently have few beliefs about the efficacy of various treatments for cancer. If I learn that I have cancer, the issue will become greatly important for me and evidence bearing on it will become all-things-considered quite important. However, the consequentiality of such evidence will not seriously change, because immediately after learning that I have cancer, I still have few beliefs on the matter. (Although if I continue inquiring about it long enough, I will acquire a sizeable body of beliefs about it such that evidence bearing on it will eventually become consequential. That is, typically reliable evidence bearing on a relatively all-things-considered important topic will become relatively consequential in the long run.) A full account of how much of one's investigative resources one ought to dedicate to scrutinizing a given piece of purported evidence would have to take all these complicated factors into account. However, we can safely bracket them in what follows. In typical examples of biased assimilation in the social psychology literature, the pieces of confirming and disconfirming evidence are comparable in other epistemic and practical respects. And, as I shall argue momentarily, closing inquiry is rooted in how the

¹⁵A Quinean confirmation-holist might hold that when any one of one's beliefs changes, it will affect one's epistemic attitudes towards all other beliefs. However, this doesn't undermine the importance of the notion of consequentiality, because it allows for changes in different beliefs to have weaker or stronger effects on one's totality of beliefs.

¹⁶See Joyce (2005) for a detailed discussion.

consequentiality of a given piece of evidence for a belief changes as a function of one's confidence in that belief, all other factors being the same.

2.2. Credence and consequentiality

Given the Key Fact, biased assimilation and closing inquiry can be shown to be grounded, rather straightforwardly, in the nature of the dependence of the expected consequentiality of evidence for H on one's confidence in H .

First consider closing inquiry. The *amount* of difference obtaining evidence, e , ought to make in one's credence in H is, by Bayes's theorem,

$$D_{H,e} = \text{df } |P(H/e) - P(H)| = P(H) \left| \frac{1}{p(H) + \frac{p(e|-H)}{p(e|H)} p(-H)} - 1 \right| \quad (1)$$

Note that $D_{H,e}$ is the absolute value of change in credence. For simplicity's sake, suppose whatever effect learning e has on one's totality of epistemic attitudes happens via its effect on one's attitude towards H (i.e., on H itself or on beliefs connected to it). Given this assumption, $D_{H,e}$ is closely linked to the consequentiality of e for an agent with a high credence in H . $D_{H,e}$ depends on two factors: $P(H)$ and $\frac{P(e|H)}{P(e|-H)}$, but only $\frac{P(e|H)}{P(e|-H)}$ is related to e . Indeed, as Howson and Urbach (2006: 97) observe, "the evidential force of e is entirely expressed by the ratio $\frac{P(e|H)}{P(e|-H)}$ known as the *Bayes factor*." (Hereafter, I will refer to $\frac{P(e|H)}{P(e|-H)}$ as the Bayes factor of e for H and denote it by ' $B_{H,e}$.')¹⁷ Two pieces of evidence with equal Bayes factors for a given hypothesis will be confirmationally equivalent for that hypothesis. Now here are the important facts. For $0 < P(H) < \frac{\sqrt{B_{H,e}-1}}{B_{H,e}-1}$, $D_{H,e}$ is a monotonically increasing function of $P(H)$ regardless of whether e confirms or disconfirms H .¹⁸ While for $\frac{\sqrt{B_{H,e}-1}}{B_{H,e}-1} < P(H) < 1$, $D_{H,e}$ is a monotonically decreasing function of $P(H)$. At $P(H) = 0$, $D_{H,e}$ is 0: evidence bearing on H has no consequentiality. As $P(H)$ becomes larger, the consequentiality of evidence for it increases. This continues until $P(H) = \frac{\sqrt{B_{H,e}-1}}{B_{H,e}-1}$ after which point, the consequentiality of e continually decreases and gets to 0 at $P(H) = 1$. If H is extremely probable (i.e., if it is sufficiently larger than $\frac{\sqrt{B_{H,e}-1}}{B_{H,e}-1}$), evidence bearing on it is inconsequential. Suppose your credence in H is 0.9 and you obtain purported evidence, e , bearing on it. If $\frac{\sqrt{B_{H,e}-1}}{B_{H,e}-1}$ is sufficiently smaller than 0.9, it might not be worth it to investigate e , because e won't make much of a difference to your credence in H , even if it is genuine evidence. In that event, you might rationally discount e ; i.e., keep your inquiry in H closed in face of e .¹⁹

¹⁷In the statistical literature, the Bayes factor is usually defined as $\frac{P(e|H_1)}{P(e|H_2)}$, which is understood as a measure of how strongly evidence e supports hypothesis H_1 compared with H_2 . My use of this concept corresponds to a special case of that definition in which $H_2 = -H_1$. Since to confirm H is to confirm it relative to $-H$, this ratio is a measure of how strongly e confirms H .

¹⁸ $\frac{\sqrt{B_{H,e}-1}}{B_{H,e}-1}$ is smaller than $\frac{1}{2}$ for confirming evidence and larger than $\frac{1}{2}$ for disconfirming evidence.

¹⁹Notice that how high $P(H)$ must be for H to be considered "highly probable" so that it is rational to keep one's investigation of H closed in face of e depends on e (or more precisely $B_{H,e}$). That is, closing

Biased assimilation is similarly grounded in the Key Fact, in only a slightly more complicated way. Here the important fact is that disconfirming evidence for a belief is always more consequential than *comparable* confirming evidence. To show this I need to first define the notion of comparability. Consider the following example. Suppose a factory makes coins that are either four times more likely to land heads or four times more likely to land tails. T is the hypothesis that a certain coin I obtained from the factory tends to land tails 80% of the time. A series of 5 throws 4 of which landed tails (s_1) is confirming evidence for T ; while a series of 5 throws 4 of which landed heads (s_2) and a series of 500 throws, 400 of which landed heads (s_3) are both disconfirming pieces of evidence for T . But clearly s_2 , and not s_3 , is intuitively comparable in force to s_1 relative to T . Earlier I mentioned one way two pieces of confirming evidence or two pieces of disconfirming evidence can be said to be comparable: when they have equal Bayes factors. A proper generalization of this idea can provide a suitable proposal for how to define comparability between pairs of confirming and disconfirming evidence. Consider two pieces of evidence, e and d , for which we have $B_{H,e} = B_{-H,d} > 1$ ($B_{H,e} = \frac{1}{B_{H,d}}$). e confirms H and d disconfirms it and the Bayes factor of e for H is equal to that of d for $-H$. Thus, e and d are in mirroring confirmational relations with respect to H . For all a ($0 \leq a \leq 1$), the confirmational effect of e for H at $P(H) = a$ is exactly equal to that of d for $-H$ at $P(-H) = a$. I think pairs of evidence that have this feature (like s_1 and s_2 in the above example) are comparable in evidential force.

Having this definition of comparable pieces of evidence, the question is: What are the relative consequentialities of comparable bodies of confirming and disconfirming evidence for H ? For $P(H) > \frac{1}{2}$, the disconfirming piece, d , is more consequential than the confirming piece, e ;²⁰ for $P(H) < \frac{1}{2}$, the opposite is the case; and at $P(H) = \frac{1}{2}$, the two are equally consequential. And $\frac{D_{H,d}}{D_{H,e}}$ is a monotonically increasing function of $P(H)$. It follows that when $P(H)$ is close to $\frac{1}{2}$, e and d are about equally consequential and it is *not* reasonable to dedicate more resources to examining d than e . However, if $P(H)$ is well above $\frac{1}{2}$, d is significantly more consequential compared with e . In that event, one ought to dedicate more resources to a critical scrutiny of d (relative to e), i.e., to exhibit biased assimilation, because if accepted as genuine evidence, accommodating d requires making more severe changes to the totality of one's epistemic attitudes.

Consider hypothesis T again. Suppose two friends of mine (in whose competence and reliability I am equally confident) tell me that they have experimented with the coin. One of them claims she has tossed it a thousand times and it landed tails 80% of the time. My other friend claims that she has done the same and the coin landed heads 80% of the time. Obviously, these pieces of evidence are comparable with respect to T . If $P(T) = \frac{1}{2}$, they are equally consequential. However, if my credence in T is meaningfully above $\frac{1}{2}$, the second report is more consequential for me. If nothing has gone wrong in the process by which it was obtained, my friend is not joking, etc., it will lead to a bigger change in my credence in T . If I have a relatively high credence in T , and I

inquiry might be rational in face of one piece of purported evidence but not another piece. This renders further support for the idea that closing inquiry is neither necessary nor sufficient for outright belief, since outright belief is not a relative notion (relativized to a piece of evidence).

²⁰The relationship between credence and outright belief is not obvious. However, plausibly whatever the nature of that relationship, a necessary condition for belief is having credence above $\frac{1}{2}$. Assuming this is true, whenever one believes p , evidence against p is more consequential than comparable evidence for p .

am willing to check the video footages, I must pay more attention to the second (disconfirming) report. This is the case where I would rationally exhibit biased assimilation. Notice, however, that if my credence in T is very high – for example, I have myself experimented with the coin a million times – I shouldn't spend any time watching either video footage (unless I am *extremely* interested in knowing the truth about T); that is, I should keep my investigation of T closed. In that event, the second report is still more consequential than the first, but neither one is particularly consequential to warrant dedicating investigative resources to it.

Before ending this section, I must make an important clarification. It might appear that my argument implicitly assumes an anti-conservative view of belief. I have argued that closing inquiry is reasonable when the evidence, even if genuine, wouldn't lead to a significant change to one's views and that biased assimilation is reasonable when and to the extent that disconfirming evidence, if genuine, results in a larger change (than confirming evidence) in one's views. Thus, it might appear that I'm making an implicit assumption to the effect that change in one's epistemic attitudes is intrinsically valuable. However, no such assumption is at work here. Rather, my argument rests on considerations of the practical (instrumental) value of higher-order evidence that bears on a piece of purported evidence. The idea is that other things being equal, inasmuch as the balance of the totality of one's evidence doesn't change significantly by the addition of a new piece of purported evidence, e , seeking higher order evidence about e (investigating e) should have low priority in the allocation of investigative resources.

3. The dynamics of belief

That a practice is licensed (or demanded) by either epistemic or practical rationality does not rule out its having adverse theoretic consequences. Probably no other epistemic principle enjoys the favorable status of consistency as a normative principle. Yet being impeccably consistent might lead to suboptimal theoretic results, because sometimes tolerating a level of inconsistency facilitates major transitions in one's views to better overall theoretic positions, which transitions are very difficult to make in a consistent manner. A famous example can be found in the debate over the implications of Newton's theory of gravity.²¹ Newton maintained that gravity was a real force of nature – not just a useful construct for saving the phenomena – and was *the cause* of many terrestrial and celestial phenomena. In the *General Scholium* of the *Principia* he wrote: “gravity really exists, acts according to the laws that we have set forth, and suffices for all the motions of the heavenly bodies and of our sea” (Newton 1999: *Principia*, vol. II: 764). Yet, the real existence of such a force (at least) appeared to contradict the conception of material causation, which was virtually universally held at the time, according to which, matter by its nature cannot act at a distance. Newton himself wrote to Bentley:

That gravity should be innate, inherent, and essential to matter, so that one body may act upon another at a distance through a vacuum without the mediation of anything else, by and through which their action and force may be conveyed

²¹The details of this important historical debate are, unsurprisingly, controversial and I cannot even remotely do them justice here. But the historical nuances are not essential to my general point. Thus, although I think what I claim here is more or less historically accurate, the reader who disagrees with my understanding of either Newton or Leibniz may suppose that I am discussing possible history that might have happened involving imaginary philosophers Newton* and Leibniz*.

from one to another, is to me so great an absurdity, that I believe no man who has in philosophical matters a competent faculty of thinking can ever fall into it. (Newton 2014: 136)

But how could a commitment to the real existence of gravity be reconciled with this understanding of material causation? This question was especially pressing given Newton's strong opposition to vortex theories, which seemed to many as the only way to find a cause for gravity that was compatible with the nature of matter. In this way, Newtonian mechanics involved attributing existence to something whose existence appeared to be ruled out by a conception of material causation that even Newton himself was not willing to abandon. Leibniz (1989) criticized Newton for engaging in "barbaric physics" and insisted that if gravity is to be accepted, it must be rendered compatible with the nature of matter (by positing certain vortices that accounted for the movements of the planets, thus rendering gravity "intelligible"). This was a much more coherent view. However, it would have had calamitous prospects for the future of science if it was adopted by everyone, since of course no satisfactory mechanistic explanation was to be found.

The strong commitment of Newton's followers to gravity was greatly beneficial for the progress of science in the long run. For it allowed them to adopt the crucial, but at the time almost unthinkable, conception of matter as a substance capable of acting at a distance. It was only after (and as a consequence of) the triumph of Newton's theory of gravitation that the accepted view of matter changed to such a substance. It took some time for the Newtonians to make this leap,²² but when they did, their theory was in better theoretic shape than Leibniz's. However, the move from Newton's own system to that of his later followers involved too big a jump that neither Newton nor his adversaries were willing to make instantly. I do not want to rule out the *possibility* of an epistemic agent who could have leaped from the pre-Newtonian conception of matter to the fully developed view of Newton's followers overnight. However, I think it is a safe bet that a natural philosopher in Newton's time was much more likely to end up accepting such overall body of beliefs (as that of later Newtonians) through a piecemeal transition, during which her theories would be partly inconsistent yet the final theoretic position she would end up with was consistent and better than any alternative she had at the outset.

This example makes it evident that a practice is not immune to having untoward theoretic consequences simply because it is rationally demanded.²³ When it comes to biased assimilation and closing inquiry, two particular effects are especially worrisome: that they seem to act as a bias in favor of one's beliefs and that they appear to increase the degree to which one's epistemic condition is path-dependent.

3.1. Path-dependence and bias

Before discussing these worries, I must clarify what I mean by the path-dependence of one's epistemic condition. Here I will largely borrow from Kelly's account of this

²²Interestingly, even Hume, who certainly admired Newton and his law of gravitation, claimed in the *Treatise* (Hume 1978) that contiguity is a conceptual element of the very idea of causation (he did not think this was incompatible with Newton's theory). This claim is absent in the first *Enquiry* (Hume 1975), probably because Hume did not want even to appear to contradict Newton.

²³Consistency is a demand of epistemic rationality. At the end of this section, I will offer another example of untoward consequences of practices rooted in considerations of practical rationality.

notion. He makes a distinction between a narrow and a broad sense of evidence. “Evidence in the narrow sense consists of relevant information about the world.” While evidence in the broader sense also includes “everything of which one is aware that makes a difference to what one is justified in believing” (Kelly 2008: 22). Hereafter, I will call evidence in the narrow sense ‘data’ and reserve the term ‘evidence’ for the broad sense. Suppose Afra receives over time an array of data ($d_1, d_2, d_3, \dots, d_n$) concerning a hypothesis, H . By reflecting on the data she might come up with various potential explanations for it. Such explanations are part of the evidence she has on H and affect which epistemic attitude(s) towards H is (are) reasonable for her to have. Kelly observes that,

Historical facts about when one acquires a given piece of evidence might very well make a *causal* difference to which body of total evidence one ultimately ends up with. One acquires a given piece of evidence at an early stage of inquiry; this might very well influence the subsequent course of inquiry in various ways, by way of making a difference to how one subsequently thinks and acts (which possibilities one considers, which routes get explored as the most promising and fruitful, and so on). And this in turn can make a difference to what evidence one ends up with. In such cases, there is an undeniable element of path-dependence. (Kelly 2008: 628)

In this picture, to the extent that one’s total evidence, and thereby the totality of one’s epistemic attitudes, are sensitive to the *manner* (order, timing, etc.) in which one acquires data, one’s epistemic condition is path-dependent.

That the epistemic conditions of agents like us are path-dependent is just inevitable. But path-dependency comes in degrees. And the lesser the better. Kelly argues that when one learns that one’s believing as one does is the result of contingent, chancy processes and one is aware of the “direction in which one’s actual total evidence is likely to be skewed,” one has reason to doubt those beliefs that are likely to have benefited from the underlying processes. I agree. Lopsided path-dependence is symptomatic of bias and is a reason to doubt the reasonableness of one’s doxastic attitude. Generally, I think learning that a practice makes one’s epistemic condition lopsidedly path-dependent has at least two normative consequences.²⁴ Upon learning this fact one ought to be less confident in *beliefs* that are likely to have benefited from the practice in question. And, learning this gives one a reason to try to avoid the *practice* to the extent possible.

3.1.1. Closing inquiry

Does either closing inquiry or biased assimilation subject one’s epistemic condition to a substantial degree of path-dependence? And does either practice work as a bias *for* one’s beliefs?

Here is why one might think the answers to both questions are affirmative with respect to closing inquiry. If one closes one’s investigation of p , one dismisses evidence for or against p unless that evidence is strong enough to justify re-opening one’s investigation. Suppose e_1 and e_2 are comparable confirming and disconfirming pieces of evidence, respectively, that bear on p . Also suppose neither evidence is strong enough to justify re-opening an investigation of p . Then because e_2 is more consequential than

²⁴In fact, I think this is the case (albeit, to a lesser extent) even if one learns that one’s epistemic condition is non-lopsidedly path-dependent.

e_i , dismissing both effectively works as a bias for p . Therefore, everything else being equal, closing inquiry acts as a bias for one's beliefs.

Closing inquiry also introduces a certain element of path-dependence to one's epistemic condition in the short run. In Afra's example above, suppose that the first r pieces of data she obtains make her confident enough about H that she closes her investigation of it in face of disconfirming data, D , and D is not strong enough to justify re-opening her investigation of H . Had she acquired D first, she might not have closed her investigation of H (her confidence in H might have been not high enough to license closing her investigation of H) and as a result, her confidence in H would have been less than it currently is, even though the data she possessed would have been identical with what she does now.

That closing inquiry does have such effects on one's epistemic condition in the short run is, I think, undeniable. Fortunately, however, this is not seriously worrisome: one closes one's investigation of p in face of evidence e in the first place because e wouldn't have made much of a difference even if one had considered it as part of one's totality of evidence. The real worry concerns the long run. If the effect of closing inquiry accumulated when evidence mounted, then a far bigger worry would have been justified. This isn't the case though (at least when closing inquiry works in isolation), because as soon as the evidence is strong, one ought to re-open one's investigation of the belief in question.

Objection: agents often won't be confronted with strong enough counterevidence to justify re-opening inquiry unless they actively seek it (which closing inquiry forbids them from doing). To use an example where this is a live worry, "many people think that the racist beliefs underlying colonialism (that 'natives' were happy, stupid and well-treated) were sustained because most Europeans not living in the colonies had only very limited evidence against these beliefs, and didn't follow up and ask for more. Of course ... they eventually re-opened inquiry when evidence mounted, but that took a while and things became very ugly in the meantime."²⁵

This objection misplaces the guilt in what was wrong in those agents' response to evidence. Closing inquiry is justified only if one has a high enough credence in a belief such that evidence concerning the belief tends not to change one's views about it. One should never close one's investigation of a belief one has very limited evidence about. Moreover, one should not mistake consequentiality for all-things-considered importance of evidence. Evidence concerning how the 'natives' were treated was relatively inconsequential (in the technical sense in which I am using it) for an agent with few beliefs about the matter. However, such an agent still had an ethical duty to know about such facts, especially if (as is assumed in this objection) what she believed was ultimately efficacious in the way those people were treated. Sometimes evidence bearing on a belief is inconsequential but, not only shouldn't one close inquiry, but one ought to actively seek evidence bearing on that belief if what one ends up believing is of great non-epistemic significance.

3.1.2. *Biased assimilation*

The adverse consequences of biased assimilation appear more worrisome (than those of closing inquiry), especially given the known causal relation it has with belief polarization. It seems that exhibiting biased assimilation can make one's epistemic condition *radically* path-dependent. Here is how. Suppose two (epistemically and practically)

²⁵This objection was raised by an anonymous reviewer for another journal.

ideally rational agents obtain different initial bodies of data that favor opposing views on a given topic, as a result of which they hold opposing views. Suppose they subsequently acquire sizeable identical bodies of mixed evidence. Since they are sensitive to demands of practical rationality, they will treat the new evidence differently. Assume they are both able to find explanations for the portion of data that is against their views that justify dismissing such data – this might not be unlikely, since they are (by exhibiting biased assimilation) dedicating more resources to critically scrutinizing that part of the data. If this happens, they will each become more confident in their beliefs. Thus, if they continue acquiring identical bodies of mixed evidence, it would seem not unlikely that they will be further polarized and will end up with radically different epistemic attitudes, even though the totality of data they have are identical except for the initial sets. If this was true, it would have been deeply troubling. Fortunately, this way of thinking about belief polarization is inaccurate.

Biased assimilation may act as a bias for one's views in two ways. First, by subjecting confirming and disconfirming data to different amounts of scrutiny, one raises the chances of the event that the disconfirming piece is discounted but the confirming piece is accepted as genuine evidence. Then the total evidence confirms the agent's initial belief because of the accepted confirming evidence. Inasmuch as this situation is a result of the unequal amount of scrutiny dedicated to different parts of the evidence – for example, in a case where if one had dedicated the same amount of critical scrutiny to confirming evidence, one would have found good reason to discount it too – biased assimilation acts as a bias for one's beliefs. However, this mechanism cannot make biased assimilation a powerful bias, because in exhibiting biased assimilation one does not carefully examine confirming evidence in the first place because and to the extent that it will not make much of a difference *even if it is genuine evidence*. The effect of biased assimilation on one's epistemic condition *through this mechanism* is very similar to closing inquiry. In the short run, it acts as a weak bias for one's beliefs but as soon as confirming evidence mounts and becomes consequential, one ought to subject it to serious scrutiny.

There is another, sometimes much more powerful, way in which biased assimilation may act as a bias for one's beliefs. Often times the fact that one has been able to find an explanation for purported anomalous data that reconciles it with one's belief is *itself* evidence for that belief. Consider the discovery of Neptune. Astronomers first detected an anomaly in Uranus's orbit, which constituted purported disconfirming data for Newtonian mechanics. Assuming the truth of Newton's theory, some auxiliary assumption about the case must have been false. In this case, what was false was the hypothesis that Uranus is the farthest planet in the solar system. By assuming the truth of Newton's laws, Leverrier and Adams were able to predict the existence of Neptune, which when confirmed became one of the most significant pieces of evidence for Newton's theory. Johann Franz Encke (the director of Berlin observatory, where Neptune was first observed) wrote to Leverrier, "your name shall henceforth be associated with the most glorious imaginable demonstration of the correctness of universal gravitation" (*Centenaire de la naissance de Le Verrier* 1911: 20, as quoted by Lequeux 2013: 34). Insofar as biased assimilation increases the chances of one's being able to turn purported anomalous data into neutral or confirming evidence (hereafter, 'friendly explanations'), it acts as a bias for one's views. And unlike the previous case, here what makes us exhibit biased assimilation – the relative inconsequentiality of confirming evidence – has nothing to do with how powerfully this practice might affect our beliefs. It is quite possible that finding friendly explanations for the 'unfriendly' data is a significant

credence-booster for our original views. In fact, the bigger the challenge purported recalcitrant data poses for a belief p , the more critical scrutiny one ought to subject such data to, and the greater impact one's ability to find friendly explanations will have on one's confidence in p . Therefore, *through this mechanism* biased assimilation might have the effect of a relatively strong bias for one's beliefs.

But this is only half of the story. Here biased assimilation may act as a bias for one's beliefs because by one's dedication of one's resources to critically examining purported recalcitrant data, one raises the likelihood of finding friendly explanations for them. However, this does not *guarantee* that one will find friendly explanations. If such explanations prove elusive, biased assimilation acts as a bias *against* one's beliefs. This is because the more one has tried to find friendly explanations, the stronger a reason against one's belief one's inability to find them will be. Thus, whether biased assimilation acts as a bias for or against one's beliefs depends on whether one is able to find plausible friendly explanations for anomalous data.

Importantly, the demands of precision and simplicity on plausible friendly explanations are sensitive to the size of anomalous data. A small body of anomalous data can be plausibly explained away with less detailed and less precise explanations than a sizeable body of such data. Here is why. Observation is subject to error. That is, even if one's theory is absolutely true, one expects to see *some* error in the observations. Therefore, small amounts of anomalous data can be explained away as "expected error" and, thus, need not be accounted for in any serious detail. When the size of disconfirming data is more than expected, one needs some other explanation for dismissing it. However, even small defects in the experimental setting can explain small deviations from the expected error. From a statistical point of view, such minor problems effectively increase the expected variance of error, which increases the acceptable error bar. However, when the amount of error exceeds a certain level, an appeal to minor defects in the experiment might no longer be a good enough explanation for the anomalous data. In that event, deeming the disconfirming evidence error requires an increasingly more detailed explanation for why the (purported) error is as large as it is. Therefore, the larger the body of disconfirming data is, the higher the precision bar for plausible friendly explanations of it will be.

The overall simplicity of one's beliefs is another factor that makes a difference in the acceptability of friendly explanations. In the short run, one is typically able to explain anomalous data away without worrying about the added complexity. However, with larger amounts of unfriendly data, considerations of simplicity might become quite forceful. If one keeps explaining the recalcitrant data away through different explanations for different pieces of data, one must make sure the addition of all those explanations do not render one's totality of beliefs too complicated or more complicated than rival overall views that are consistent with one's data. Here one's explanations might act like Lakatosian protective belts that are not reasonable to add after a certain point.

In sum, when one dedicates more resources to scrutinizing anomalous data, one increases the chances of finding plausible friendly explanations for it, *if there are such explanations to be found*; but one also makes it more difficult to reasonably stick to one's belief in the event that one is unable to find such explanations. (If there are plausible friendly explanations for the counterevidence, biased assimilation effectively helps the epistemic standing of one's beliefs, but this is not worrisome, because in such cases the theory has earned this epistemic standing.) Hence, given the different standards for friendly explanations in the short and the long run, and given the fact that biased assimilation works as a bias for or against one's beliefs

depending on whether one is able to find friendly explanations, biased assimilation is much more likely to act as a bias for one's beliefs in the short run, but these effects are mitigated to a fair degree, if inquiry is carried out long enough.

Therefore, contrary to initial appearances, biased assimilation does not have major worrisome consequences for one's epistemic conditions in the long run, at least when considered in isolation.

3.2. Forgetting evidence

The dynamics of one's beliefs become far more worrisome when these conservative practices are combined with epistemic traits, such as forgetting evidence, that might hinder or slow the accumulation of counterevidence even in the long run. Forgetting evidence is not in principle problematic for agents with limited memory, like us. As Harman has observed,

There is a limit to what one can remember, a limit to the number of things one can put to long-term storage, and a limit to what one can retrieve. It is important to save room for important things and not clutter one's mind with a lot of unimportant matters. ... one should try to restore in long-term memory only the key matters that one will later need to recall. When one reaches a significant conclusion from one's other beliefs, one needs to remember the conclusion but does not normally need to remember all the intermediate steps involved in reaching that conclusion. Indeed, one should not try to remember those intermediate steps. (Harman 1986: 42)

Considerations of practical rationality forbid us from trying to remember every single unimportant detail or step taken in reaching a conclusion. Likewise, and by virtually the same reasoning, one must not try to remember every single piece of evidence for one's beliefs. Thus, not only is it natural for one to forget the evidence, at times it is a demand of rationality not to actively try to remember it, in order to save room in one's memory for more important matters. However, inasmuch as this might hinder the growth of the size of counterevidence, it tends to make us *strongly* biased in favor of our already-held beliefs and make our epistemic conditions highly sensitive to initial conditions.

The combination of biased assimilation or closing inquiry with forgetting evidence might have seriously worrisome effects on one's epistemic condition, if the rate at which one acquires data is small and/or the rate at which one forgets data is large. In such cases, one ought to (i) put more effort to remembering *discounted counterevidence* than one otherwise would; and (ii) be less confident of beliefs that are likely to have benefited from this combination.

References

- Christensen D.** (2010). 'Higher-Order Evidence.' *Philosophy and Phenomenological Research* 81, 185–215.
- Feldman R.** (2000). 'The Ethics of Belief.' *Philosophy and Phenomenological Research* 50, 667–95.
- Feldman R.** (2009). 'Evidentialism, Higher-Order Evidence, and Disagreement.' *Episteme* 6, 294–312.
- Foley R.** (1987). *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.
- Foley R.** (1993). *Working Without a Net*. New York, NY: Oxford University Press.
- Frankish K.** (2009). 'Partial Belief and Flat-Out Belief.' In F. Huber and C. Schmidt-Petri (eds), *Degrees of Belief*, pp. 75–93. New York, NY: Springer.

- Harman G.** (1986). *Change In View: Principles of Reasoned Revision*. Cambridge, MA: MIT Press.
- Howson C. and Urbach P.** (2006). *Scientific Reasoning: The Bayesian Approach*, 3rd edn. La Salle, IL: Open Court.
- Hume D.** (1975). *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*, 3rd edn. Oxford: Clarendon Press.
- Hume D.** (1978). *A Treatise of Human Nature*, 2nd edn. Oxford: Clarendon Press.
- Joyce J.** (2005). 'How Probabilities Reflect Evidence.' *Philosophical Perspectives* **19**, 153–78.
- Kelly T.** (2003). 'Epistemic Rationality as Instrumental Rationality: A Critique.' *Philosophy and Phenomenological Research* **66**(3), 612–40.
- Kelly T.** (2008). 'Disagreement, Dogmatism, and Belief Polarization.' *Journal of Philosophy* **105**(10, Special Issue), 611–33.
- Kelly T.** (2010). 'Peer Disagreement and Higher Order Evidence.' In A. Goldman and D. Whitcomb (eds), *Social Epistemology: Essential Readings*, pp. 183–217. Oxford: Oxford University Press.
- Lasonen-Aarnio M.** (2014). 'Higher-Order Evidence and the Limits of Defeat.' *Philosophy and Phenomenological Research* **88**, 314–45.
- Leibniz G.W.** (1989). *Philosophical Essays*. (R. Ariew and D. Garber, eds and transl.). Indianapolis, IN: Hackett.
- Leitgeb H.** (2014). 'The Stability Theory of Belief.' *Philosophical Review* **123**(2), 131–71.
- Lequex J.** (2013). *Le Verrier-Magnificent and Detestable Astronomer*. New York, NY: Springer.
- Lord C., Ross L. and Lepper M.** (1979). 'Biased Assimilation and Attitude Polarization: The effect of Prior Theories on Subsequently Considered Evidence.' *Journal of Personality and Social Psychology* **27**, 2098–109.
- Newton I.** (1999). *Principia: Mathematical Principles of Natural Philosophy*. (B. Cohen and A. Miller Whitman, eds and transl.). Berkeley, CA: University of California Press.
- Newton I.** (2014). *Philosophical Writings*. (A. Janiak, ed.). Cambridge: Cambridge University Press.
- Shaffer M.J.** (2019). 'Explaining Evidence Denial as Motivated Pragmatically Rational Epistemic Irrationality.' *Metaphilosophy* **50**(4), 563–79.
- White R.** (2005). 'Epistemic Permissiveness.' *Philosophical Perspectives* **19**, 445–59.

Alireza Fatollahi is a postgraduate fellow at Princeton University and will soon join the philosophy department at Bilkent University as an assistant professor. He got his PhD and MA in philosophy from Princeton and University of Houston, respectively. Before that, he got his BSc in physics from Sharif University of Technology.