

8-2018

Epistemic dimensions of gaslighting: peer-disagreement, self-trust, and epistemic injustice

Andrew D. Spear
Grand Valley State University, speara@gvsu.edu

Follow this and additional works at: https://scholarworks.gvsu.edu/oapsf_articles




Part of the [Epistemology Commons](#)

ScholarWorks Citation

Spear, Andrew D., "Epistemic dimensions of gaslighting: peer-disagreement, self-trust, and epistemic injustice" (2018). *Funded Articles*. 127.

https://scholarworks.gvsu.edu/oapsf_articles/127

This Article is brought to you for free and open access by the Open Access Publishing Support Fund at ScholarWorks@GVSU. It has been accepted for inclusion in Funded Articles by an authorized administrator of ScholarWorks@GVSU. For more information, please contact scholarworks@gvsu.edu.



Epistemic dimensions of gaslighting: peer-disagreement, self-trust, and epistemic injustice

Andrew D. Spear

To cite this article: Andrew D. Spear (2019): Epistemic dimensions of gaslighting: peer-disagreement, self-trust, and epistemic injustice, Inquiry, DOI: [10.1080/0020174X.2019.1610051](https://doi.org/10.1080/0020174X.2019.1610051)

To link to this article: <https://doi.org/10.1080/0020174X.2019.1610051>



© 2019 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 25 Apr 2019.



Submit your article to this journal [↗](#)



Article views: 575



View related articles [↗](#)



View Crossmark data [↗](#)

Epistemic dimensions of gaslighting: peer-disagreement, self-trust, and epistemic injustice

Andrew D. Spear

Philosophy Department, Grand Valley State University, Allendale, MI, USA

ABSTRACT

Miranda Fricker has characterized *epistemic injustice* as “a kind of injustice in which someone is wronged specifically in her capacity as a knower” (2007, *Epistemic injustice: Power & the ethics of knowing*. Oxford: Oxford University Press, 20). *Gaslighting*, where one agent seeks to gain control over another by undermining the other’s conception of herself as an independent locus of judgment and deliberation, would thus seem to be a paradigm example. Yet, in the most thorough analysis of gaslighting to date (Abramson, K. 2014. “Turning up the lights on gaslighting.” *Philosophical Perspectives* 28, *Ethics*: 1–30), the idea that gaslighting has crucial epistemic dimensions is rather roundly rejected on grounds that gaslighting works by means of a strategy of assertion and manipulation that is not properly understood in epistemic terms. I argue that Abramson’s focus on the gaslighter and on the moral wrongness of his actions leads her to downplay ways in which gaslighters nevertheless deploy genuinely epistemic strategies, and to devote less attention to the standpoint and reasoning processes of the victim, for whom the experience of gaslighting has substantial and essential epistemic features. Taking these features into account reveals that all gaslighting has epistemic dimensions and helps to clarify what resistance to gaslighting might look like.

ARTICLE HISTORY Received 1 August 2018; Accepted 1 February 2019

KEYWORDS Gaslighting; epistemic self-Trust; self-Trust; peer-Disagreement; epistemic injustice

In this essay I argue that gaslighting has essential epistemic dimensions. Gaslighters can offer their victims straightforward (if fabricated) epistemic reasons as part of the gaslighting process, and every gaslighter makes at least a tacit claim to a position of epistemic superiority relative to his victim, one that plays a role in the victim’s experience of gaslighting. From the victim’s standpoint, gaslighting calls into question not just her moral or psychological, but also her epistemic self-trust: her conception

CONTACT Andrew D. Spear  speara@gvsu.edu

© 2019 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

of herself as an independent locus of experience, thought, and judgment.¹ Because the victim trusts both herself and her gaslighter, his challenge to her self-trust generates, from her standpoint, a situation of epistemic peer-disagreement where what she and her gaslighter disagree about is specifically her own epistemic standing, and so her entitlement to epistemic self-trust. In this connection I discuss the conditions under which it is reasonable for someone to accept defeat for their own epistemic self-trust (and so to conciliate in the disagreement with the gaslighter), and argue that these conditions are particularly pertinent to understanding the situation of the victim of gaslighting as well as what well-grounded resistance to gaslighting might look like. In stressing the epistemic dimensions of gaslighting I differ from the analysis of the issue offered by Kate Abramson (2014), even while following her account closely in other respects.

Section 1 introduces the phenomenon of gaslighting by means of examples, and discusses Abramson's analysis of gaslighting. Section 2 outlines Abramson's reasons for downplaying the role of epistemic factors in gaslighting, while Section 3 begins a response to these by discussing epistemic dimensions of gaslighting on the side of the gaslighter. Sections 4 and 5 focus on the epistemic dimensions of gaslighting as experienced by the victim, while section 6 considers the implications of the account offered here for resistance to gaslighting.

1. The gaslighting phenomenon and Abramson's analysis

The term 'gaslighting' derives from Patrick Hamilton's 1938 play *Gas Light* and from the 1944 film of the same name starring Charles Boyer and Ingrid Bergman. In the early 1980s the term became the center of some discussion in psychoanalysis (Calef and Weinshel 1981), and it has since been picked up in the realms of self-help (Stern 2007) and political commentary (Carpenter 2018). While no treatment of the topic is likely to accommodate all of these different uses, Kate Abramson has provided a thorough philosophical analysis of gaslighting and of its moral wrongness (Abramson 2014). While accurate and illuminating in many ways, Abramson's

¹Here and throughout I use 'she', 'her', etc. to refer to the victim of gaslighting and 'he', 'his', etc. to refer to the perpetrator. This usage helps to avoid pronoun-confusion throughout, and fits with the situations in key examples of gaslighting (the films *Gas Light* and *Pat and Mike*) and with the feminist slant of Abramson's 2014 analysis of the phenomenon. In adopting this usage, I do not mean to take a stand concerning whether gaslighting is particularly or predominantly perpetrated by men against women. I think that the agent and patient of gaslighting can be mixed in any number of ways in terms of gender (men to women, men to men, women to men, etc.) and other social relations. Who gaslights who predominantly or the most, in terms of social category, is an empirical question that I expect depends rather heavily on social and cultural factors.

account focuses on the role of manipulation in gaslighting while downplaying epistemic factors such as testimony, evidence, and reasons. I think, however, that these latter are equally essential to a complete understanding of the phenomenon. In this section I will discuss two paradigm examples of gaslighting and introduce Abramson's account.

1.1. Two central examples of gaslighting

Abramson begins her discussion by introducing a variety of cases that are plausibly considered gaslighting. I'll consider two of these here in order to frame the discussion.

First, there is the case presented in the 1944 film *Gas Light*. In the film, Paula (Ingrid Bergman) is the victim of gaslighting on the part of her husband Gregory (Charles Boyer). Unbeknownst to herself, Paula is in possession of priceless jewels that belonged to her aunt, a famous opera singer who was mysteriously murdered when Paula was a young girl. Gregory, who is in fact the murderer of Paula's aunt and single-mindedly bent on gaining possession of the jewels, romances and marries Paula in Italy, then convinces her to return to London to live in her aunt's old house. Gregory is confident the jewels are in the attic of the house and 'goes out on business often' in order to secretly go to the attic by a back staircase to search for the jewels. His use of the gaslights in the attic causes the gaslight in the rest of the house to go dim. Paula sees this, and often hears mysterious footsteps in the attic. Gregory not only assures her that she is imagining things, but from the beginning suggests to her that she is overtired, and that her memory is not working properly. He subtly arranges things so that it appears, to Paula and others around her, that she has been stealing, hiding, or moving things around in the house. Of course, when he confronts Paula with these things, she has no recollection of having done them, which only confirms Gregory's 'suspicion' of her faulty memory and imagining of things. As a result, Paula's confidence in her own judgment and mental faculties dramatically deteriorates. Gregory's goal is to drive her to the point where she believes that she is mad so that he can get her out of the way and have the house to himself to search for the jewels. The 'gaslighting' in the film is Gregory's attempt to convince Paula not to trust her own judgment and faculties, and so not to trust herself.

The second case, one that plays a central role in Abramson's analysis, is also from a film. This time from the 1952 Film *Pat and Mike*. In the film Pat is an aspiring female golfer, while her fiancée Collier wants (tacitly at least) for her to give up her golf career so that they can get married and she

can assume the wifely duties of household and children. After a close tournament loss in which Collier's less than supportive approach has played a clear (to the viewer) but not explicit role, Pat and Collier have the following exchange:

Collier	How about looking on the bright side of this for instance? Take this—As long as your job's out of the way, move the date up, tie the old knot? I think you've done enough, worked long enough, don't you?
Pat (distressed)	oh, too much (looks down)
Collier (interrupting Pat)	After all, what you trying to prove, who you trying to lick?
Pat (determinedly, upset)	Myself. (pounds fist in air)
Collier	Just the kid who'll do it (Puts his leg up and looks at her dubiously and patronizingly)
Pat	Collier, do you sort of, I don't think you mean to, but do you think of me as just the little woman?
Collier	That's right, and myself as a little man. (Squeezes her shoulders like a small child).
Pat (quite distressed)	Right now, now I feel like a sort of flop that you're rescuing. I'm flummoxed, that's what I am. Maybe we ought to wait until I don't feel so carved up, so nobody.
Collier	Why don't you just let me take charge!
Pat (fatigued)	I have to be in charge of myself.
Collier	Oh what's the good of that, I mean after all?
Pat	I have to have time to think it over.
Collier	Well, just make sure you don't think it under. It's a nice long ride, just take your time. (Opens newspaper to end conversation). (The screenplay for <i>Pat and Mike</i> , quoted in Abramson 2014, 7)

In this case, and in many of the cases Abramson focuses on, the gaslighting is more subtle than in the case of Paula and Gregory, but has a similar goal. What Collier wants is for Pat to stop resisting his vision of their future together, even if that vision essentially involves that she give up on projects that are important to her. He aims to achieve this both by calling into question, however subtly, her ability to manage things for herself ('I think you've done enough, worked long enough, don't you? ... Just the kid who'll do it ...'), and also by tacitly communicating to her the possibility of his emotional withdrawal and so of her loss of the relationship with him if she doesn't see things his way ('Well, just make sure you don't think it under' followed by the abrupt ending of the conversation). As in *Gas Light*, so here, we have

the gaslighter attempting to undermine his victim's confidence in herself, and so to break down her resistance and get her to endorse his plans, decisions, and view of reality rather than her own.

1.2. Abramson on gaslighting

On the basis of the foregoing and similar examples, Abramson argues that gaslighters are individuals who cannot tolerate even the possibility of disagreement with or criticism of their way of viewing things, at least not from certain individuals (friends, loved ones, romantic partners ...), and that the purpose of gaslighting is not only to neutralize particular criticisms that such individuals might lodge, but to neutralize *the very possibility* of criticism by undermining the victim's conception of herself as an autonomous locus of thought, judgement, and action. Even when the strategy is subtle and manipulative, what the gaslighter really wants, and so what gaslighting centrally involves, is to undermine his victim's capacity to criticize or respond independently of him. Abramson's analysis of gaslighting is that,

... he [the gaslighter] aims to destroy the possibility of disagreement by so radically undermining another person that she has nowhere left to stand from which to disagree, no standpoint from which her words might constitute genuine disagreement. (2014, 10)

Thus,

... the *paradigm* case of gaslighting...[is] ... one in which the gaslighter *wholeheartedly, constantly and consistently* aims at the destruction [of] his or her target's standing to issue challenges ... (2014, 11)

On Abramson's view, gaslighting essentially involves manipulation. It involves (i) issuing a demand to the victim (ultimately, that she accept the gaslighter's assessment of herself and 'see things his way') where (ii) the victim's 'motive for assent' to the demand comes in the form of implicit or explicit 'manipulative threats' on the part of the gaslighter. Thus, Collier's demand to Pat is that she accept that she is not seeing the situation properly or need not worry about seeing the situation properly since, after all, *he* can; and that further, if she fails to see the situation properly then he may end the relationship (2014, 15). While Abramson does not stress the point, her analysis makes it clear that gaslighter and victim will typically stand in a significant relationship of trust or authority, such that the gaslighter is able to leverage specific interests of his victim either in him, in their relationship, or in her social or economic situation that can motivate her compliance.

2. Abramson against an epistemic reading of gaslighting

Abramson is throughout critical of the view that there is a significant epistemic dimension to gaslighting. She points out that the gaslighter is not primarily (or at all) '... trying to get a person to rethink her reactions, or to see another perspective' (2014, 13), and that many of the typical retorts of the gaslighter take the form of directives or proclamations: 'don't be paranoid' and 'that's crazy' are blunt assertions, not invitations to further discussion. Further, Abramson suggests that if the gaslighter is challenged by his victim, he does not respond by engaging any evidence she might provide, but rather by broadening or intensifying his declarative attack on his victim's capacities, motives, and limitations, finally resorting to tacit or explicit manipulation if the victim persists in her challenges. Further on, in the heart of her discussion of gaslighting as manipulation, Abramson says,

... he [the gaslighter] isn't in the first instance claiming for himself a epistemic authority (I see this rightly, you don't) ... what he's doing is issuing a demand that one see things his way ... this isn't a case of, for instance, testimonial credence (i.e. the gaslighter isn't asking his/her target to take it on testimony that it's true that "that's crazy"). If that were the scenario, there'd be no explanation for the gaslighter's use of manipulative threats (implicit or explicit). It's the explicit or implicit manipulative threats ... that give the target anything like motive for assent. (2014, 15)

Similarly, Abramson is critical of the idea that gaslighting is productively understood as what Miranda Fricker has called 'testimonial injustice'. According to Fricker, such injustice occurs when a speaker's word suffers a loss in credibility for the hearer due to *identity prejudice* (Fricker 2007, 28). For example, if a white jury member does not give the testimony of an African American witness the credence that it deserves in a courtroom due to the fact that the juror is prejudiced against African Americans. Abramson responds,

To suppose that in gaslighting, the primary issue is about credibility assessments is, I think, to focus in the wrong place. It's to lose sight of the fact that an important part of what's going on is that the gaslighter is trying to turn a situation that might involve credibility assessments into a situation in which credibility assessments are not at issue, because there is no credibility to be assessed. (2014, 17)

While I am in agreement with much of Abramson's analysis of gaslighting, on the issue of credence and the epistemic dimensions of gaslighting I think that she is too quick to be dismissive.

3. Epistemic dimensions of the gaslighter's gaslighting²

I agree that a necessary condition for gaslighting is that the gaslighter have the goal of '... destroying his or her target's standing to issue challenges ...' (2014, 11). Further, and here I think Abramson would agree, what it means to 'destroy' another person's standing to issue challenges, in the context of gaslighting, is specifically to bring it about that the other person *views herself* as deficient or completely incompetent concerning her ability to understand, interpret situations, think, and choose for herself. It would be possible to 'destroy' another's standing to issue challenges by simply killing them, or administering mind-controlling drugs, or undermining her credibility with others in relevant social contexts, but the gaslighter wants more than this: he is specifically out to achieve a change of attitude in his victim towards herself.

If this is the goal of gaslighting, however, then the following seems at least possible: the gaslighter's motivation for gaslighting could be the simple desire to gain total or near total control over the victim, while the strategy for achieving this end could be to leverage the victim's trust in him as a credible peer or authority and to provide her with a selective, carefully constructed, or outright manipulated set of evidence in order to convince her that she is indeed unable to handle all or most cognitive and deliberative functions on her own. On this model, the victim's acquiescence would be a consequence of her (from a third-person omniscient standpoint, clearly misplaced) *trust* in the gaslighter, and of the particular array of evidence or arguments (to whatever degree actually fabricated) that the gaslighter has provided. Abramson suggests that if this is what the gaslighter were doing, then there would be no explanation for the use of 'manipulative threats' by the gaslighter, and no account of the victim's 'motive for assenting'. Yet manipulative threats could certainly play a role in clouding the ability of the victim to properly assess the gaslighter's credibility and the quality of the reasons that he offers, and Abramson does not explain why straightforward epistemic reasons couldn't provide the victim with a motive for assent, especially since the 'assent' at issue is specifically supposed to be a change in belief about herself, not merely a change in behavior.

To develop this point, in typical cases of gaslighting the gaslighter wants the victim to accept his views about two things. First and foremost, he is attempting to convince her that she should not trust herself: that she

²Versions of the arguments of Sections 3 and 4 first appeared in a briefer form in (Spear 2018).

is 'crazy' or 'oversensitive', that her grasp of the situation is too defective for her to make reliable judgments about it, or that her judgments themselves can't be trusted. Second, there is typically also a specific further claim that the victim is supposed to accept: that she has been moving things around in the apartment and then forgetting, or that Pat should give up her career and marry Collier. The gaslighter thus, arguably, has a twofold project. There is the ongoing project of getting his victim to accept that she is, in general, not competent to make certain (or most) kinds of judgments and so should not trust herself, and then there is the appeal to the results of this project as a sort of premise to compel compliance or ward off critique from the victim on particular occasions. On Abramson's view, gaslighting is only happening if the gaslighter attempts to motivate his victim to accept these attitudes toward herself and what he wants by means of manipulation. Yet, as already suggested, it seems clear that gaslighting at least *can* be about the epistemic status of the gaslighter and about giving the victim primarily epistemic reasons to distrust her own experience and judgment. This is precisely the scenario of Gregory and Paula in the 1944 film. It is because Paula *trusts* Gregory as a reliable judge of how things are, and because he then systematically manufactures plausible evidence that her faculties are untrustworthy, that she is brought to the point where she is no longer confident in her own judgments. This doesn't preclude his use of more manipulative appeals to Paula's emotions and insecurities, but even when the reasons he gives her are purely (fabricated) epistemic ones, I see no reason to say that what he is doing isn't gaslighting.

The point just made about Paula's *trust* in Gregory is crucial, however, for it indicates a broader respect in which *all* gaslighting involves issues of epistemic status and trust. Whether the gaslighter challenges his victim's self-trust primarily manipulatively, by playing upon her emotional commitments or insecurities, primarily epistemically (by offering reasons or relying on his perceived credibility), or by means of some mix of these two strategies (probably the typical case), his gaslighting efforts can gain purchase only if the victim trusts him, and an essential part of this trust must, I argue, be epistemic. He maintains, after all, that she is incapable of seeing and so responding to the situation properly, and thus that she should trust him rather than her own grasp of what is going on. However, to say to someone that they fail to see a situation correctly, whether calmly and quietly, or abruptly and unpleasantly (as the gaslighter frequently does: 'that's crazy' or 'that's just what *you* would think') is to tacitly claim to see that same

situation at least somewhat better. Even if he doesn't say it in just these terms (and even if he is not sincere), by making the assertions that he does concerning his victim's grasp of the situation, the gaslighter is, contra Abramson, by this very fact laying claim to a special or privileged epistemic status relative to his victim. Another way to see this, I think, is that were the victim to seriously doubt the epistemic status (the credibility) of her gaslighter, this fact would itself render his gaslighting project unworkable. If the victim interprets the gaslighter's abrupt and aggressive behavior as indicative of some (ultimately epistemic) failure to properly *see* the situation on *his* part, rather than as arising from a failing of her own, then he may still be able to manipulate her behavior, but he will not be successful in actually *changing her own assessment of herself*, in gaslighting her.

Thus, insincere though he typically will be, I think the gaslighter's assertions do involve claims to epistemic status relative to his victim, at least tacit ones. I do not know if this is so 'in the first instance' as Abramson says it is not, but I do think it is an essential element of all gaslighting behavior. Indeed, part of the reason I think this is precisely because credibility assessments of this sort are relevant to the victim's 'motive to assent'. While I agree with Abramson that the gaslighter will typically also be manipulating his victim, I think the victim only has motivation to assent *to change her beliefs about her own grasp of the situation* if she takes the gaslighter to have an epistemically credible grasp of the situation. The gaslighter could convince her to accept this assessment of himself either in the explicitly epistemic fashion of Gregory in *Gas Light* or by deploying manipulative threats, as does Collier in *Pat and Mike*, but in both cases the credibility assessment itself plays a crucial part in the gaslighting process. Finally, while I agree also that the gaslighter is trying to '... turn a situation that might involve credibility assessments into a situation in which credibility assessments are not at issue, because there is no credibility to be assessed' (Abramson 2014, 17), so long as this means specifically getting the victim to *change her mind* about her assessment of *her own* credibility, her assessment of *his* credibility is an essential element of the process (a point I will develop in more detail in the next section). Finally, whether the gaslighter is specifically committing testimonial injustice against his victim, his attempt to undermine her own assessment of herself as a credible epistemic agent surely is an attempt to wrong her 'specifically in her capacity as a knower', and so is an epistemic injustice in Miranda Fricker's sense of this term (Fricker 2007, 20).

Thus, while Abramson's view of gaslighting as manipulation based on affective and prudential vulnerabilities of the victim captures many important aspects of gaslighting and helps to draw out the nature of the wrong it involves, her insistence that issues of epistemic reasons, credibility assessments, and testimony play little or no role in gaslighting seems mistaken. Even when viewed solely from the standpoint of the gaslighter, gaslighting *can* have an epistemic dimension insofar as the gaslighter's strategy may be overtly epistemic (as in the case of Gregory in *Gas Light*), while all gaslighting has an essential epistemic dimension insofar as the overt assertion that someone else does not grasp a situation properly always involves the tacit assertion that the speaker (the gaslighter in this case) does see the situation at least somewhat better, and so should be trusted. I will develop this point about trust on the side of the victim in the next section.

4. The victim of gaslighting: epistemic self-trust and the epistemology of disagreement

For the victim of gaslighting, the fundamental issue she confronts when experiencing gaslighting has to do with self-trust. The gaslighter's goal is to convince his victim that her thoughts, perceptions, memories, judgments, and evaluations of the situation are so unreliable that she should put little or no credence in their deliverances, and should instead see things as he sees them. Self-trust is, roughly, the tacit or explicit belief that a relevant subset of one's cognitive faculties (e.g. one's perceptual faculties, reasoning, memory, etc.) are in general aimed at the production of true beliefs (accurate inferences, etc.), that these faculties are functioning properly, and that one's assessments of their deliverances are in general and on the whole correct (veridical perceptions and accurate inferences would matter little if the subject nevertheless put these together in random or incoherent ways to form beliefs about situations and about the world). For my purposes here two claims about epistemic self-trust are important. The first is that it is default rational to have self-trust concerning one's own cognitive faculties and judgments. Subjects don't need to be in possession of special reasons or arguments in order to be justified in trusting themselves. The second is that self-trust is nevertheless defeasible. While subjects do not need special reasons to trust themselves, there are types of reason a subject could confront that might require them to suspend judgment in or even give up their self-trust.

Linda Zagzebski has recently presented a view of self-trust as *basic*, *natural*, and *rational*, but also *defeasible*. According to Zagzebski, there is a natural pre-reflective desire to have true beliefs and a natural belief that this desire is in general satisfiable, where these two things combined 'imply that there is also a pre-reflective trust in the suitability of one's cognitive faculties for arriving at truth' (2012, 36). While this self-trust is pre-reflective, Zagzebski argues that *reflective* self-trust arises as a result of the confrontation with various skeptical possibilities, and of the realization that any attempt to justify belief in the trustworthiness of one's cognitive faculties and standpoint will ultimately be circular. There is no non-circular way to show or demonstrate to one's self that one's cognitive faculties are reliable.

Further, Zagzebski understands rationality in a particular way as 'doing what we naturally do better'. On her view, we naturally desire true beliefs, and part of this desire involves using our reason to minimize or eliminate cognitive dissonance when our beliefs seem to be inconsistent either with each other or with new evidence provided by our faculties. Zagzebski argues, in effect, that the most rational thing for the reflective self to do after the confrontation with skepticism and the realization of the epistemic circularity involved in any attempt at bootstrapping epistemic self-trust, is to reflectively double-down on the natural self-trust of the pre-reflective self. Self-trust is thus natural, unavoidable if one is going to pursue truth at all, and rational by Zagzebski's understanding of rational, yet it is defeasible: '... if we regularly had inconsistent memories or unstable perceptions, we would know that something was amiss with these faculties; they could not be trusted' (2012, 41). In what follows, I will rely on Zagzebski's account of self-trust in particular in order to more closely analyze the epistemic nature of the situation of the victim of gaslighting.

If the goal of the gaslighter is to undermine his victim's self-trust in a global way (so as to undermine her ability to meaningfully criticize him and/or to gain control over her), then the victim's experience of gaslighting must be one in which the default and even reflective rationality of her self-trust is meaningfully and quite globally challenged. So, what might such a challenge look like? The answer, I think, is that the gaslighter begins the process of gaslighting by introducing cognitive dissonance, often quite emotionally charged cognitive dissonance, into the relationship with his victim, dissonance that specifically requires the victim to decide between her own way of viewing things and that of the gaslighter, and the gaslighter then works directly or indirectly to ensure that the victim resolves the dissonance in his favor, specifically by downgrading

her conception of herself as a locus of independent thought and judgment relative to him. More specifically, from the standpoint of the victim of gaslighting, the experience is that of epistemic peer-disagreement of the sort dealt with in recent discussions in epistemology, but with a twist.

Two individuals are epistemic peers if they are approximately equal in their informedness, reasoning abilities, freedom from bias, and other cognitive performance and judgement-relevant factors. The central question in recent discussions of epistemic peer disagreement in epistemology has been that of what the rational thing to do is when two individuals who take themselves to be epistemic peers disagree about some proposition, whether something simple such as the correct splitting of a shared dinner check, or something more controversial such as a political or moral question. Does the mere fact of such disagreement, keeping all other evidence constant, constitute grounds for the individuals involved in the disagreement to adjust their credence in their respective beliefs downward (the 'conciliationist' position) or not (the 'steadfast' position) (see e.g. Christensen 2009)? The idea of an epistemic peer naturally suggests the idea of both an *epistemic superior* (or epistemic authority) and an *epistemic inferior*. It is not difficult to imagine someone attempting to resolve apparent peer-disagreement by down-grading their interlocutor, by revoking her status as epistemic peer and so the claim of her opposing belief to serious rational consideration. A slight variation on this would be to attempt to gain concessions from or control over one's interlocutor by *convincing her* that she herself is not one's epistemic peer in some significant respect. I take this move to be a central element in gas-lighting: the gas-lighter attempts to gain concessions from his victim by reducing her, by her own lights, to an epistemic inferior or dependent.

More to the point, and this is the 'twist' on ordinary peer-disagreement that I suggested above, what the gaslighter does is make the question of the reliability of his victim's cognitive faculties (her ability to grasp, interpret, and correctly judge the situation) *itself* the proposition about which they disagree. This places the victim in a very difficult epistemic situation because she must weigh her trust in her own epistemic agency against her trust in the gaslighter as an epistemic peer, bearing in mind that his calling into question of her cognitive faculties arguably provides at least some defeating evidence for her epistemic self-trust, evidence that she cannot reject in a non-circular fashion (Christenson 2009, 760). However, the debate between the conciliationist and steadfast positions gets resolved, it would be very difficult to deny that there are *some* cases where it is rational for a subject to take the word of a trusted

friend, acquaintance, or authority as grounds that her cognitive faculties are malfunctioning. For instance, a close friend might inform one that one has been misremembering things often lately, and that this may well be a sign of approaching Alzheimer's. From the victim's standpoint, I suggest, gaslighting often mimics such cases. The victim finds herself confronted with a situation where she must weigh her epistemic self-trust (her trust in her abilities to perceive, reason, and form beliefs adequately) against her trust in the testimony and assertions of an epistemic peer or authority, where it is her epistemic agency itself that is being called into question.

Thus, from the standpoint of the victim of gaslighting, gaslighting is a profoundly epistemic phenomenon having to do both with questions of epistemic self-trust and with the issue of epistemic peer- or even authority-disagreement. Whether the gaslighter proceeds according to a strategy of assertion and manipulation (as in the example of Collier in *Pat and Mike*) or according to a more overtly epistemic strategy of manipulatively providing false but misleading evidence to his victim (as in the example of Gregory from *Gas Light*), what the victim finds herself confronted with is a challenge to her epistemic self-trust being lodged by a person she views as a typically sincere epistemic peer or authority. If she revokes her epistemic self-trust, then she can maintain her view of the gaslighter as a typically sincere epistemic peer or authority. If she revokes her view of the gaslighter as a typically sincere epistemic peer or authority, then she can retain her epistemic self-trust. Yet, the cognitive dissonance created by the gaslighter means that she can't do both, even as she may have very strong emotional and prudential reasons (of the sort discussed insightfully and at length by Abramson 2014, 18–23) to concede to the gaslighter. It is only by recognizing the essentially epistemic nature of the victim's situation, however, that we can fully appreciate the quandary she faces and the options available to her, and this is an essential part of understanding both gaslighting and what successful resistance to it might look like. An important question remains, however, and this is the question of when or under what conditions the victim is justified in resisting or in capitulating to gaslighting?

5. The victim of gaslighting: reasons and self-trust-undermining reasons

I take it that reasons can cause beliefs and motivate actions, and reasons for both belief and for action can be better or worse, more or less justified:

acting for good reasons makes an agent count, all things being equal, as rational. The class of reasons is not, however, homogenous, as there are arguably prudential reasons, moral reasons, epistemic reasons, aesthetic reasons, and etc. The final question that I want to consider here is that of when, or on the basis of what reasons, it is justified for the victim of gaslighting to resist her gaslighter (when is it reasonable to treat the gaslighter's behavior as a reason to downgrade trust in him rather than in one's self?) and when it is not. The question at issue is thus that of when or under what conditions it would be reasonable for a subject to believe that her own cognitive agency had been compromised and thus that she should no longer trust herself to perceive, think, and judge independently. My suggestion here is that belief in one's own epistemic agency, one's epistemic self-trust, is special. While there may be beliefs that it is acceptable to hold for pragmatic, prudential, or moral reasons *even though* one lacks good epistemic reasons for them, and while there may be beliefs that it is acceptable to abandon for similar reasons and in spite of lacking good epistemic defeaters, belief in one's own epistemic agency is not among these beliefs. What this means is that gaslighting has yet another epistemic dimension. Namely, the victim who capitulates to the gaslighter and cedes her self-trust for any reasons other than legitimate epistemic ones, when not doing so is within her power, is violating what I suggest is a basic epistemic norm.

5.1. Epistemic reasons and prudential reasons

That there are different things we care about and so different kinds of reasons for action is a relatively common assumption in recent and even not so recent philosophy. Call an *epistemic reason* a reason that makes beliefs based on it more likely to be true. I take epistemic reasons to be the kinds of reasons that, when possessed to a sufficient degree by a subject, typically confer epistemic justification on her beliefs (BonJour 2003, Ch. 1; Bergmann 2006, Ch. 1). Similarly, and perhaps most importantly here, I take epistemic reasons to be the kinds of reasons that comprise defeaters for belief. A subject *S* has a defeater *D* for belief in the proposition *P* just in case *D* makes it all things considered likely that *P* is false (a 'rebutting defeater') or *D* undermines *S*'s original reason for believing that *P* (an 'undercutting defeater') (Pollock 1986; Plantinga 1993, Ch. 12). A subject who believes that the French Revolution began in 1778 (sparked off, as he believes, by the untimely death of Jean-Jacques Rousseau) and then is told by a renowned historian and expert on the

French Revolution that it in fact began in 1789 receives a rebutting defeater for his belief. He should now think it is *false* that the French Revolution began in 1778. A subject who believes that the French Revolution began in 1778 because she read it on a particular website, and then receives information that the website in question is riddled with errors and so untrustworthy, has an undercutting defeater for her belief. She has reason to be agnostic about the truth of ‘the French Revolution began in 1778’ (assuming the website was her only source of information about this). Epistemic reasons, including defeaters, provide subjects with reasons for thinking that a proposition is true (or false) independently of what they wish or desire to be the case. For example, a forlorn liberal in the United States who followed the 2016 Election results November eighth and ninth of that year on a series of credible news sources had strong epistemic reasons for believing that Donald Trump would be the next President of the United States, regardless of what they desired or judged was best. Given this understanding of epistemic reason, a person can be said to be *epistemically rational* if they proportion their beliefs to the epistemic reasons or evidence that they possess.

By contrast with an epistemic reason, call a *prudential reason* a reason that some individual has for desiring or wanting a state of affairs to be the case. A subject has a prudential reason when something would be ‘good for’ her: it would enhance her actual or perceived well-being, at least to some extent (Crisp 2008, Sec. 1; Taylor 2013). I take it that prudential reasons are first and foremost reasons for acting. If Fred has left his laundry on the line and does not want it to get (or remain) wet, he has a prudential reason to go out and take it down ahead of the coming thunder storm. Taking the laundry down is good for him because he won’t have to expend the effort to dry it a second time. If Sally aspires to be a talented concert pianist, she has prudential reasons to practice every day and to seek admission to a high-quality conservatory for training. This would be good for her. I take it that there can also be prudential reasons for holding beliefs. If a subject has a prudential reason for holding a belief, this just means that it would be ‘good’ or ‘better’ for her to hold the belief than not. All things being equal, it is arguably good for most subjects to believe that they are going to live past tomorrow, that their spouses are faithful, that their conceptions of themselves as persons are approximately accurate to how they actually are, and that the rainy weather will not be perpetual. A subject is then prudentially rational if they act (or form beliefs) in a way that is all things considered ‘good for’ them.

While holding true beliefs will, I think, often be prudentially rational such that epistemic rationality and prudential rationality will align, there is no reason this must always be the case and I suspect it often is not.³ William Clifford's famous dictum that 'it is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence' (Clifford 1877/1886/1877/1886) suggests that the only reasons governing the rationality of belief are epistemic reasons. William James, in part responding to Clifford, argued that we should not be so restrictive, and that our so-called 'passional nature' might be entitled to base doxastic commitments, at least under certain conditions involving high stakes values such as religious beliefs or the meaning of life (what James calls 'genuine options') even where *no* evidence is present: thus beliefs without epistemic reasons *can* be rational since they are based on other kinds of reasons (James 1896/2010/1896/2010). More mundane cases seem to offer a similar lesson. In some contexts, such as sports, believing that one's abilities are greater than they are, even in the face of evidence to the contrary, may actually have a positive influence on performance. Similarly, in the context of health and disease, it may be the case that believing against the epistemically available odds that one will recover has positive benefits, such as making one's illness more bearable or increasing one's chances of recovery (Feldman 2003, 43–44). Further, trust is arguably an important value in interpersonal relationships, yet it may be harmful or even contradictory to insist that one's trust, e.g. in one's spouse's faithfulness, be fully or adequately supported by epistemic reasons or evidence (McLeod 2015). Further, it may be that the underdetermination of theory by evidence in science means that much scientific theory choice is, in part, an expression of social, evaluative, or aesthetic reasons on the part of scientists, and so not purely epistemic (Kuhn 1977), and this may even extend to the majority of our extra-scientific epistemic practices and judgments (Jaggar 1989). Thus, there are many ways in which epistemic reasons and prudential reasons might come apart or come into conflict for a given subject.

³Resolving the issue is in any case beyond the scope of this essay. A strict Aristotelian view of well-being may well be one on which prudential rationality and epistemic rationality strictly align, while a subjective-preference-satisfaction view of well-being is, I think, likely to lead to divergence between prudential rationality and epistemic rationality. Whatever account of well-being is accepted, phenomena such as confabulation and self-deception seem to point rather clearly to ways in which epistemic rationality and prudential rationality of belief can come apart (Mele 2002; Hirstein 2009; Sullivan-Bissett 2015).

5.2. A basic epistemic norm concerning the defeat-conditions for epistemic self-trust

For my purposes here, I am willing to concede all of the foregoing and will not even make much effort to distinguish amongst the different points just reviewed. What I want to argue is that, however the foregoing points get worked out, one's epistemic agency and so epistemic self-trust are *distinctive* in relation to all of the cases discussed so far in that it is not possible to form *any* rational beliefs for any reasons whatsoever if one does not first believe in one's basic ability to evaluate reasons and form beliefs in response to them. If an agent accepts that her cognitive faculties, perceptions, judgments, and overall grasp of her own situation are all substantially deficient, then she is no longer in a position to endorse as rational, under any understanding of 'rational', the judgements she arrives at, and so is no longer able to meaningfully go forward as an agent. Not only is she not in a position to trust her purely epistemic judgements about the way the world is or about how likely a given proposition is to be true, but she will be unable to meaningfully trust her prudential and evaluative judgments as well insofar as these too are products of the cognitive faculties that she has accepted are broadly malfunctioning or untrustworthy. Yet this is exactly what the gaslighter wants his victim to do, to cease believing that she is an autonomous locus of thought, deliberation, and belief. If he is successful in this project, then his victim will no longer be even a potential source of criticism or resistance to him, and he will have nearly total control over her.

Given the inquiry-ending and agency-undermining consequences of accepting such a belief about one's self, my argument here is that the following is a basic epistemic norm concerning defeat for one's own epistemic self-trust:

It is rational for an agent to accept that her epistemic agency is globally undermined *only* for epistemic reasons.

In particular, an agent should not accept defeat for her self-trust for prudential reasons, such as desiring to avoid negative personal or professional consequences, or the desire to maintain an intimate relationship or friendship. Because comprehensively ceding self-trust undermines one's very ability to meaningfully be concerned with prudential reasons to begin with, no merely prudential reason should count, from an agent's own standpoint, as a reason sufficient for giving it up.

Epistemic self-trust is, as discussed above, default and reflectively rational to have, but it is subject to defeat. That defeat, I am arguing

here, must however take the form of legitimate epistemic defeat, not merely involve an agent giving up on her self-trust for moral or prudential reasons. None of this implies that it can never be rational to concede that one's epistemic agency is compromised. It is typical for patients suffering from certain kinds of neurological or cognitive disorders not to be fully aware of this fact about themselves (Hirstein 2009, introduction). However, if such a person is confronted with diagnoses (e.g. of the early stages of Alzheimer's disease) from credentialed doctors that are also corroborated by trusted friends or family members who are able to bring to the agent's attention particular examples of their own recent behavior that further corroborate the diagnosis, I think that such an individual is being confronted with significant reasons to doubt her epistemic self-trust. Even here, the distrust should be carefully proportioned to the extent of the cognitive malfunction for which the agent has some evidence, but it would be epistemically arrogant not to take such considerations into account. It is true that a gaslighter who uses what I have called an 'epistemic strategy' of gaslighting, may well try to mimic precisely such an evidential situation (as does Gregory for Paula in the film *Gas Light*), but this just highlights the point that our epistemic agency, while fundamental, is also fragile such that we are all, in principle, susceptible to gaslighting.⁴

An additional point to be noted here is that the epistemic norm I am proposing applies primarily to a subject's 'global' or 'total' epistemic self-trust. There is nothing epistemically problematic in preferring the deliverances of someone else's cognitive faculties over one's own on particular topics or in well-circumscribed areas (Zagzebski 2012, Ch. 1). Perhaps someone else's vision or mathematical abilities are better than my own, and so I defer to them on these matters. Similarly, it is no threat to epistemic self-trust to rely on credible authorities in areas where one is not an expert (for most of us, the majority of areas outside of commonsense and a few other select domains). None of the foregoing are cases where the norm I have proposed is clearly violated. Further, though I cannot think of many examples, I would be willing to countenance the possibility that a subject might permissibly give up epistemic self-trust in some well-

⁴There are two additional issues here. The first is that an individual who suffers from cognitive defects, vices, or malfunction may, for this very reason, be unable to appreciate and rationally respond to even the most significant evidence that this is the case about her (see, e.g. Kornblith 1998). The second is that there is a kind of paradox involved in an agent who rationally judges and so accepts that her own cognitive faculties are not trustworthy; after all, shouldn't this judgement itself now be untrustworthy? I think that both of these points are important and raise further issues that need to be addressed, but plead limitations of space here.

demarcated area for merely prudential reasons (in the absence of any epistemic reasons). Conceding to someone else's memory of just what happened at the party last night (to the point of genuinely believing their account is correct) simply because it is easier than the painful argument that would otherwise follow *may* not be so problematic.⁵ Regardless, the norm I am proposing here is intended to apply to self-trust understood in a rather global way, as a break-down in one's conception of one's ability to grasp, process, and judge concerning facts in non-specialized ordinary domains in which large differences of ability in judgment and knowledge are atypical or, in any case, where there is no reason to think that they exist between one's self and the individual or individuals who are calling one's basic grasp of things into question. It is in such cases, the typical cases where gaslighting occurs, where accepting defeat for one's self-trust violates the epistemic norm at issue.

6. Epistemic self-trust, gaslighting, responsibility, and resistance

What can the foregoing account tell us about resistance to gaslighting? On what grounds might victims of gaslighting, at least in principle, push back? First, I have argued that the victim ought not concede her epistemic self-trust for anything but well-credentialed epistemic reasons. While the victim's personal investment in the gaslighter and the gaslighter's deployment of manipulative threats may present obstacles to its consistent application, recognition of this principle itself provides significant grounds for resistance. It implies that gaslighting behavior should typically count as defeating evidence for the credibility of the person engaging in it, excepting only the case where the gaslighter is so skilled at providing his victim with (fabricated) epistemic reasons for doubting her epistemic self-trust that an objective assessment of the situation, from her own point of view (considering the epistemic reasons actually available to her), requires accepting defeat.

Second, even in such a case, capitulation may not be required, and also for reasons accessible from the subject's point of view. Robert Pasnau, addressing the literature on epistemic peer disagreement, has recently argued that self-trust is something intrinsically valuable that stands over

⁵I am, frankly, inclined to think that it *is* problematic if the subject goes so far as to form *the belief* for merely prudential reasons. It is one thing to decide (for prudential reasons) not to argue about the matter anymore. It is another to actually change one's mind. But I don't need to settle this question in order to make the point I want to make here about the global norm as it relates to gaslighting.

against epistemic rationality (which is also intrinsically valuable on his view) (Pasnau 2015). Pasnau argues that if epistemic rationality were the only thing agents were concerned about, then some form of conciliationism would seem like the best response to epistemic peer disagreement across the board. Yet, he argues that conciliationism is problematic because, at least in certain important cases, it runs up against self-trust, which requires a certain coherence between one's beliefs and one's evidential perspective. Even if my epistemic peer is, from an impartial perspective, equally rational, informed, and etc., from my own perspective it may be the case that things still seem to me, in light of all available evidence, to be a certain way. Because self-trust is independently valuable, Pasnau argues, it is at least permissible in such cases to continue to give my own assessment of the matter and so my own beliefs greater weight (and thus to lean in the direction of the steadfast approach to peer disagreement).

While Pasnau's discussion does not imply that it is reasonable for agents to give greater weight to their own beliefs no matter what contrary evidence they confront, it does give self-trust particularly great weight in determining what it is reasonable to believe in cases of peer-disagreement. This is particularly relevant for the victim of gaslighting, insofar as her situation involves both an epistemic peer-disagreement and a challenge to her self-trust. If Pasnau's arguments are correct, then it may be the case not only that an agent should not concede her epistemic agency for anything other than significant epistemic reasons, but also that even in the face of such reasons, the value of preserving epistemic self-trust makes it reasonable for her to persist in her own beliefs, indeed in maintaining trust in herself, even in the face of a certain amount of evidence to the contrary. Recognizing the way in which the victim's experience of gaslighting involves epistemic disagreement thus opens up, via considerations such as Pasnau's, avenues for meaningful and well-grounded resistance.

If the epistemic norm that I have identified and the potentially defeating value of self-trust argued for by Pasnau both suggest that resistance to gaslighting will be the rational and justified response in many cases, then this raises the question of responsibility when victims capitulate. The account I have offered here does imply that at least some victims of gaslighting who capitulate to their gaslighters are to some degree responsible for this capitulation. This will be so when two conditions are met. First, the victim must in fact be in possession of sufficient reasons concerning herself, her gaslighter, and their situation

that the epistemically reasonable thing to do when confronted with his gaslighting behavior is to retain her own epistemic self-trust, probably while downgrading her trust in the gaslighter. Second, the victim must be psychologically able, at least in principle, to correctly recognize and process the reasons available to her. In saying this, it may seem that I am investing the victim of gaslighting with more power and ability to respond to the situation than she may typically have. *Ought implies can*, after all, so saying what a victim ought to do concerning beliefs about her own self-trust seems moot if the reality of her situation is such that she is simply overwhelmed and so unable to meaningfully do other than she does in capitulating. In this connection, Abramson discusses a significant list of ‘tools’ of the gaslighter whereby he attempts to manipulate his victim, such as the victim’s own love for him or the confidence deficit that many women suffer from as a result of social conditioning and expectations (Abramson 2014, 19–23). I do not mean to downplay such factors or the difficulty of the situation faced by the victim of gaslighting. Nor do I deny that there are cases of gaslighting where rational resistance of the sort I have been describing here is simply not an option for the victim. Indeed, I think such cases are possible and likely quite frequent. However, such cases will be cases where at least one of the two conditions I have identified is not satisfied, and so where the victim who capitulates is *not* responsible (or is not *fully* responsible, depending on the case). Further, whether the victim who capitulates bears some type of responsibility or not, this *in no way* diminishes the responsibility of the gaslighter for the wrong that he is committing against her.

The gaslighter’s goal is to completely neutralize his victim’s conception of herself as an independent locus of thought and criticism. The fact that, with the help of additional manipulative strategies or a massive power imbalance, this goal may be achieved more efficiently and with less resistance in some cases does not change the general dynamic at work. Perhaps social conditions or interpersonal relations are such that some (or even many) agents start out with epistemic self-trust so fragile that it takes little (or even nothing?) to successfully gaslight them. This doesn’t show that gaslighting is not an epistemic phenomenon or that agents are never responsible for capitulating to gaslighting, but rather that social conditions are such that some (and perhaps many) agents are systematically prevented from developing or having the epistemic self-trust that they are rightfully entitled to. In this regard I hope that my extension of the analysis of gaslighting to its epistemic dimensions might function both as a guide for what successful resistance to gaslighting might look like (maintaining

self-trust even in the face of attacks on it by someone one takes to be a credible peer), and also to more serious consideration of the social conditions and types of character that we should seek to cultivate and allow to flourish in order to make its use and success less common and less likely.

7. Conclusion: gaslighting and its epistemic dimensions

In the foregoing I have extended Abramson's analysis of gaslighting to include a fundamental epistemic dimension. I take myself to agree with the majority of Abramson's analysis of the characteristic goals, methods, and moral wrongness of the gaslighter, while disagreeing with her only in her downplaying of the phenomenon's essential epistemic dimensions. Yet this disagreement is significant. The account on offer here suggests that *Gaslighting* involves (i) the attempt by one person, the gaslighter, to undermine his victim's self-trust: her conception of herself as an autonomous locus of experience, thought, and judgment; that the gaslighter's (ii) motivation is to gain or maintain control of his victim specifically by means of getting her to internalize a view of herself that both neutralizes her ability to criticize him and ensures her consent to his way of viewing things (specifically with regard to issues relevant to the relationship, perhaps in general); that the gaslighter (iii) pursues this goal by means of a strategy of manipulation, fabrication, and deception that (iv) specifically relies upon his victim's trust in him as a peer or authority in some relevant sense.⁶ The most distinctive feature of gaslighting is that it is not enough for the gaslighter simply to control his victim or have things go his way: it is essential to him that the victim herself actually *come to agree with him*. Thus, gaslighting is distinct from mere silencing, from creating an environment where *everyone else* believes the victim is wrong, and also from creating a situation where the victim has no choice but to acquiesce, even while not agreeing.

So understood, there are four epistemic dimensions to gaslighting. The first is that the gaslighter himself may use a strategy of providing false or fabricated evidence to his victim in order to undermine her self-trust, as Gregory does to Paula in the film *Gas Light*. Second, in asserting that his victim lacks a clear grasp of the situation, the gaslighter is always at

⁶Gaslighting is not limited to intimate partners in a relationship. As characterized here, gaslighting is possible in many interpersonal contexts involving trust or authority, such as employee-employer relationships and relationships amongst peers of various sorts (friends, co-workers, fellow students, etc.).

least tacitly asserting that he has a better grasp, and this is a kind of claim to epistemic status or authority. Third, from the standpoint of the victim, gaslighting always places her in a situation of epistemic peer disagreement with her gaslighter, where the point of disagreement is specifically whether her own cognitive faculties are reliable or not, and she must weigh her trust in the gaslighter and associated evidence for this against her own self-trust. Fourth, and finally, self-trust is a fundamental epistemic value and the condition of all other meaningful epistemic agency, thus an agent should not accept defeat for her self-trust on anything but epistemic grounds. The agent who capitulates to gaslighting for merely prudential or moral reasons thus violates a fundamental epistemic norm. While there may be agents for whom no possibility of resistance to gaslighting exists, the foregoing account still helps to make clear what meaningful resistance, where possible, might look like, and also points to the project of articulating the social and interpersonal conditions under which legitimate self-trust might flourish, thus minimizing the occurrence or possibility of successful gaslighting.

Acknowledgements

For suggestions and conversation about gaslighting in all its forms I am especially grateful to Katherine Tullmann, Stephanie Adair, and Jeffrey Byrnes. An earlier version of this paper received helpful comments and criticism from all of the participants at the Grand Valley State Philosophy Summer Research Group, and from participants at the 2018 meeting of the European Epistemology Network. The final version of this essay is far better than it otherwise would have been thanks to comments and suggestions from an anonymous referee.

Disclosure statement

No potential conflict of interest was reported by the author.

References

- Abramson, K. 2014. "Turning Up the Lights on Gaslighting." *Philosophical Perspectives* 28, *Ethics*: 1–30.
- Bergmann, M. 2006. *Justification Without Awareness*. Oxford: Oxford University Press.
- BonJour, L. 2003. *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*. New York: Wiley-Blackwell.
- Calef, V., and E. M. Weinshel. 1981. "Some Clinical Consequences of Introjection: Gaslighting." *Psychoanalytic Quarterly* v 50 (1): 44–66.
- Carpenter, A. 2018. *Gaslighting America*. New York: Harper Collins.

- Christenson, D. 2009. "Disagreement as Evidence: The Epistemology of Controversy." *Philosophy Compass* 4/5: 756–767.
- Clifford, W. K. 1877/1886. "The Ethics of Belief." In *William K. Clifford: Lectures and Essays*, edited by L. Stephen and F. Pollock, 177–212. London: Macmillan and Co.
- Crisp, R. 2008. "Well-Being." *The Stanford Encyclopedia of Philosophy*. Accessed November 9, 2018. <http://plato.stanford.edu/archives/win2008/entri=es/well-being/>.
- Feldman, R. 2003. *Epistemology*. Upper Saddle River: Pearson.
- Fricker, M. 2007. *Epistemic Injustice: Power & the Ethics of Knowing*. Oxford: Oxford University Press.
- Hirstein, W. 2009. *Confabulation: Views from Neuroscience, Psychiatry, Psychology, and Philosophy*. Oxford: Oxford University Press.
- Jaggar, A. 1989. "Love and Knowledge: Emotion in Feminist Epistemology." *Inquiry: An Interdisciplinary Journal of Philosophy* 322: 151–176.
- James, W. 1896/2010. "The Will to Believe." In *The Will to Believe and Other Essays in Popular Philosophy*, 13–45. New York: The Floating Press.
- Kornblith, H. 1998. "What is it Like to be me?" *Australasian Journal of Philosophy* 76: 48–60.
- Kuhn, T. S. 1977. "Objectivity, Value Judgment, and Theory Choice." In *The Essential Tension*, 320–339. Chicago: University of Chicago Press.
- McLeod, C. 2015. "Trust." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Stanford: The Metaphysics Research Lab, Center for the Study of Language and Information. <https://plato.stanford.edu/archives/fall2015/entries/trust/>.
- Mele, A. 2002. *Self-deception Unmasked*. Princeton: Princeton University Press.
- Pasnau, R. 2015. "Disagreement and the Value of Self-Trust." *Philosophical Studies* 172: 2315–2339.
- Plantinga, A. 1993. *Warrant and Proper Function*. New York: Oxford University Press.
- Pollock, J. 1986. *Contemporary Theories of Knowledge*. Savage, MD: Rowman and Littlefield.
- Spear, A. 2018. "Gaslighting, Confabulation, and Epistemic Innocence." *Topoi*. doi:10.1007/s11245-018-9611-z.
- Stern, R. 2007. *The Gaslight Effect*. New York: Random House.
- Sullivan-Bissett, E. 2015. "Implicit Bias, Confabulation, and Epistemic Innocence." *Consciousness and Cognition* 33: 548–560.
- Taylor, T. E. 2013. "Well-Being and Prudential Value." *Philosophy and Public Policy Quarterly* 31 (2): 10–17.
- Zagzebski, L. 2012. *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief*. Oxford: Oxford University Press.