



Three Puzzles with *Ad Hominem* Arguments

ABSTRACT: *The ad hominem appears to be the simplest fallacy form—one criticizes speakers instead of their statements or arguments. It is regularly taken to be a fallacy of irrelevance, in that who is speaking does not bear on the truth of what is said. But three puzzles attend this analysis. (1) Given that the fallacy is simple and seemingly obvious, how could it be effective in practice? (2) Are there not cases when who is speaking is relevant? How do we sort those cases from those where it is irrelevant? (3) Isn't there another level to the ad hominem, one where we observe it, know it is a bad argumentative move, and make inferences about the argumentative circumstances and arguers in light of it? Accusing another of committing the ad hominem, on this line of thought, has broader implications about the reasoner and their reasons. This article is an attempt to tell a coherent story of the ad hominem that makes sense of these three puzzles and shows how the observed preponderance of the argument form should trouble us as reasoners.*

KEYWORDS: *Ad hominem*, fallacy, argumentation, meta-argumentation

I.

The *ad hominem* is arguably the most famous fallacy. Both to commit and to complain about. Our proclivity as arguers to make arguments personal, and to take disputes personally, is very strong. Moreover, there are temptations to make inferences about arguers who go *ad hominem*, too. It is clear, in one sense, that critical thinking teachers and argumentation theorists should smile in appreciation when complaints about the *ad hominem* show up in popular discussions of critical dialogue. Witness:

Barack Obama (former U.S. President): “We just don’t fling out *ad hominem* attacks like that, because it doesn’t help inform the American people.” (*Remarks by President Obama and Prime Minister Hailemariam Desalegn of Ethiopia in Joint Press Conference, 2015*)

Ted Cruz (U.S. Senator): “He didn’t discuss any of the substance. Instead, he just attacked and engaged in *ad hominem* attacks.” (*Sen. Cruz Slams Dr. Fauci on ‘Hannity’: ‘The Most Dangerous Bureaucrat in the History of the Country’ | U.S. Senator Ted Cruz of Texas, 2021*)

Jennifer Rubin (*Washington Post* Contributor): “In each of these cases, their *ad hominem* attacks denigrate the debate and distract from the real point.” (2021)



Ann Coulter (U.S. political commentator): "... *ad hominem* attack is the liberal's idea of political debate" (Coulter, 2002)

There are some important take-aways here. First, it seems that the vocabulary of informal logic is in common parlance, and that people are rightly complaining about and criticizing bad arguments. And further, it's clear that everyone in the conversation knows what is happening. That's good news, right? Informal logicians should celebrate: *hooray, we're relevant!*

The trouble is that after our celebrations, it becomes clear that we still have work to do on explaining precisely what *ad hominem* arguments are and what all the consternation about them means. There are, I believe, three interrelated puzzles with the *ad hominem* that emerge as we establish a tighter focus on the phenomenon. They are, in the form of questions:

How could insulting someone be an effective argumentative strategy in the first place? It, as a fallacy, is just so obviously irrelevant. That is, there seems to be so much *ad hominem* argument out there, but how might it move anyone? Call this the *effectiveness puzzle*.

Aren't there cases where *ad hominem* argument is actually appropriate or relevant? That is, even though *ad hominem* is a fallacy of relevance, aren't there non-fallacious instances where who is speaking is relevant, and if so, how do we sort them from the irrelevant? Call this the *relevance puzzle*.

What's the significance of all of our complaints about the *ad hominem*? It is clear that the fallacy is widespread, and it is equally clear that we are sensitive to the fallacy as a fallacy. In this reflective mode, it seems that we are making inferences about the overall state of arguments given around us, particularly those with whom we disagree. Are these good inferences? It seems there's a lot of that fallacy out there, and nobody's happy about it. That's weird, right? Moreover, people seem to be making all sorts of inferences about the state of argument generally from the fact that there are lots of *ad hominems* given. Are these good inferences? Call this the *meta-argumentation puzzle*.

My plan here is to lay out the puzzles and form replies to them, constructing a criterion for relevance, an explanatory scheme for how and why *ad hominem* fallacies can be effective at (appearing to be) successful argument, and what follows from the fact that we feel we are awash in arguments that amount to character assassination of the arguers. I think a worrying lesson emerges for us, one that arises as evidence about our meta-argumentative reflection; so, a meta-meta-argumentative reflection on how we are likely not processing the evidence as well as we should, even when we are unhappy about all the fallacious arguments.

2.

One place to start is with the personalizing element of so much of our argumentative and informative communication. Take the simple phenomenon of inferring *Some S are not P* from being informed that *Some S are P*. Despite what introductory logic students might think, one doesn't *just mean* the other, and some easy work with Venn diagrams and lessons about the bottom part of the Square of Opposition go a long way in explaining why. (The fallacious inference has the not-so-famous name of *illicit subcontraries*—if only that could make its way into popular parlance as *ad hominem* has!) And now the question emerges as to *why* the illusion of the formal fallacy's validity exists. Why does it *seem* like they *just mean* each other? The answer is that it is not an inference about the *information* in the claims, but about the *informers* making the claims. The move of mind is a *speaker-regarding inference*. Some speaker A attests that "Some S are P," and a competent listener B may reasonably infer that A also is communicating that some are not.¹ The competence here arises from an assessment of the speaker from some background assumptions about communication, and the listener reasons: *were A to know that 'all S are P,' she would have said so, and consequently, since she said only some are, she is also communicating that some are not*. The big lesson from all of this is that we not only reason about the things we discuss with each other, we reason about each other. We assess each other as informers, reasoners, and worthwhile conversational partners. All that talking is energy-intensive, and it is even more so when we are arguing. So we owe it to each other to make it more efficient, and we are regularly assessing how we are doing at getting at the truth, arguing, and making it all work smoothly. And the next lesson, returning to the *Some S are P* inference, is that it is all automatic and feels as though we are just getting to the items at issue. Speaker-regarding inferences come back around to the issue, and in this they feel like they've been about the issue all along.

Ad hominem arguments are cases of attacking or criticizing a person advancing a viewpoint in order to attack the viewpoint. In this regard, the notion of a speaker-regarding inference is a powerful explanatory tool for identifying what's happening with the *ad hominem* – we reason about the informer as a piece of automatic interpersonal information processing, and then spit out an inference about the information. In this case, the basic form of *ad hominem* arguments is:

Speaker A attests that p

A is of objectionable character

Therefore, p is not acceptable

¹ The Gricean line adds that we both know this offline inference is being made, so in saying "Some S are P" to communicate that some are not, we must rely on the speaker knowing that the hearer knows that the speaker knows the hearer will make the inference. Speaker-regarding inferences are, then, not only automatic and counter-factual, but they are also meta-cognitive.

The trouble with the inference is that a proposition's truth or falsity is not changed by the character of the person accepting or asserting it.² Moreover, a terrible person can still have excellent reasons for asserting a claim and commending it to others. Consequently, the error of the *ad hominem* is classified as a *fallacy of relevance*—we make an inference about a claim's quality from something irrelevant to that claim (namely, the claimant). The objectionable character of a claimant, in a good deal of the literature on the fallacy and its textbook treatment, can come in a variety of forms, traced by further distinctions. There is the *abusive* form of the argument, which amounts to a simple insult or a moral objection to the ethics of the claimant. There is the *circumstantial* form, which is an observation of conflict of interest or incapacity with the evidence for the claimant. And there is the *tu quoque*, which is an observation of an inconsistency either of the claimant's claims and other claims, or their claims and their actions.

Now that we have stated the inferential frame for the *ad hominem*, the orienting puzzles can be posed in ways that can be slightly clearer:

The effectiveness puzzle: Given the irrelevance of *ad hominem* attacks, how could an argument like this change any minds? How does this reasoning ever seem good?

The relevance puzzle: Aren't there times when who is speaking is relevant? How do we sort the relevant from irrelevant cases?

The meta-argumentation puzzle: We seem to both object to and relish *ad hominem* arguments, and that seems strange on its face. Further, we make inferences about the broader argumentative situation when we see *ad hominem* arguments prevail, particularly about how informed our interlocutors are or how good their overall case is. So, how do those inferences work?

Fallacy theory, by my lights, entails many objectives. They include being able to explain what's wrong with the fallacious arguments, explain why we fall for them, identify how to re-rail the derailed conversation, and explore what the significance of the error being widespread is. Finally, it's worth asking whether all this knowledge creates a virtuous loop of reflection and action or a vicious one, a cycle of rationalization borne of our newly gained reflective capacities. I'll argue, to close, that the loop is a mixed bag, but this is the fate of fallible and reflective creatures like ourselves.

3.

Fallacious arguments are bad arguments that seem good, and a core objective of fallacy theory is to address how these bad arguments might seem to be good. The effectiveness puzzle is simply that it seems, on its face, implausible that *ad hominem*

² The breadth of agreement in the literature on this point is significant. See: (Hamblin, 1970; Hinman, 1982; McMurtry, 1986; Brinton, 1995; Walton, 1998; Woods, 2007; Walton, Reed and Macagno, 2008; Yap, 2013, 2015; Aikin and Talisse, 2019; Wisley, 2019; Hundleby, 2023)

arguments, particularly the egregiously fallacious ones, could ever even seem to be good arguments. How could insulting an interlocutor's haircut make him change his mind about economics? How could making fun of an arguer's terrible choice in clothes make them think they are wrong about vaccine effectiveness? How could accusing me of being a drunk make me rethink my political opinions? (I'll hasten to add that I drink so much precisely because *I know a good bit about politics!*)

A distinction between the kinds of audience addressed will be useful. There is a difference between the *ad hominem* being given in the *second person* (YOU) and the *third* (HE, SHE, THEY). Consider the following basic forms:

Second person: You are contemptible, so you are wrong.

Third person: They are contemptible, so they are wrong.

The issue here is what kind of dialogue model we are using to make sense of the argumentative exchange and whose mind one is out to change. On a *dyadic* model, we have two discussants trying to resolve their disagreement with reasons facing each other. On a *polyadic* model, we have at least one extra participant, an *onlooking audience* to the exchange, and participants are trying to move *their* views on the issue (see Macagno, 2013; Aikin and Casey, 2022b; Lewiński and Aakhus, 2014; Lewiński, 2019; Aikin and Talisse, 2019; Aakhus and Lewiński, 2023). The effectiveness issue looks very different once we've introduced this distinction, since the main problem was to make sense of how insulting someone might change their mind.

Let's start, then, with the second person version of the *ad hominem*, since with it, the effectiveness challenge is still in high relief. There are two ways to make sense of how the argument could be effective in the dyadic relation, but we will need to change our approach on what the form of effectiveness is. The first explanation is to hold that the *ad hominem* is not actually an inferentially pregnant contribution to the dyadic argument, but rather it is a reason given to the interlocutor to disengage with the dialogue. Consider the fact that argument is time- and energy-intensive, and if one of the interlocutors drives the cost of exchange up by being insulting, we may just avoid the exchange altogether and beg off it when we get some evidence of the fact that it'll be heavy weather to continue. So, the *ad hominem*, instead of being an argument at all, is still a *reason*—it's a reason for one of the interlocutors to disengage from the exchange. And notice how this strategy has an effective and convenient result for the one giving the *ad hominem*: the opponent, once things got heated, retreated and so has forfeited on the issue. A default meta-argumentative and interpretive view is that claims put forward without objection have tacit agreement, and *ad hominem* arguments have driven the costs of objection high enough for one's view to have the default status. Think of how reacting to people who challenge 'common sense' commitments as though they are ridiculous has not only a silencing effect on those with critical questions but bolsters the view that the commonsense views are unproblematically correct.³

³ Those committed to this indirectly inferential reading of the *ad hominem* are (van Eemeren and Grootendorst, 1995a, 1995b; Woods, 2007; Budzynska and Witek, 2014; Hitchcock, 2017; van Eemeren and Garssen, 2023)

Consider this first approach to the effectiveness puzzle to be analogous to that of the *ad baculum*, threatening others in the midst of a critical dialogue. For sure, the promise of a beating does not change others' opinions, but it will change their behavior—they will find themselves saying different things or at least remaining silent when asked if there are any objections. In this regard, the interlocutor addressed with the *ad baculum* argument is not one really to be convinced at all, but rather to be given a reason to no longer make trouble for this discussion (Casey, 2022). We might call this function, more broadly, *suppression of counter-arguments*, the tactic of making it prohibitively costly for others to play active opponent roles in arguments. With the *ad baculum*, it is with threats of violence or other punitive consequence, and with *ad hominem*, it is with active insults and implied threats of more attacks. Again, the tactic, when successfully deployed, does not yield *conviction* in targets for abuse, but it does yield *compliance*.

The second approach to the second-personal address with *ad hominem* is that it is a form of argument with the interlocutor, but its approach is to induce a kind of self-doubt. A strategic insult may make a person's identity salient, and so induce a form of stereotype threat-relevant decrease in performance. Or the insult is the front edge of a campaign of gaslighting an interlocutor into not believing their own reasons (see Yap, 2015; Aikin and Casey, 2022b; Hundleby, 2023). So, calling an interlocutor a 'snowflake' might make them second-guess their judgments about harms done to them or a marginalized group. Or calling another arguer a 'boomer' may be a way of making the interlocutor no longer trust their experience in a new and quickly developing domain. Importantly, this strategy has its effect when it is consistently applied, usually as a matter of cultural (or sub-cultural) norms invoking identity for undercutting epistemic standing. If one's interlocutor has internalized these norms (that is, they see themselves as the gaslighters see them), then all one must do is invoke the identity (perhaps in an accusatory manner) and then let the work be done internally on the opponent's side to undercut their rational self-confidence.

So, in the second-person, *ad hominem* arguments are designed to either chase an interlocutor out of the dialectical space and then claim victory,⁴ or to activate some picture the opposition may have of themselves so that they do not have the confidence in their own argumentative performances.

In the third person, *ad hominem* arguments work on an image of the *negative ethos* of the target (see Brinton, 1985, 1995; Walton, Reed and Macagno, 2008; Budzynska and Witek, 2014). In contrast, the basic form of *positive ethotic* argument is that one presents oneself in a fashion that highlights one's expertise, good judgment, and cultural competence that increases an audience's identification—they like you and trust your judgment, so they are more likely to find what you say credible. These cases, alternately, come in negative profile—one may come across as a kind of *negative authority* on the matter, in that the less an audience likes you, the more likely it is that they want to you fail—in this case, at argument. Perhaps you have met someone you find to have *negative charisma*. Everything they do is just

⁴ A good deal of this approach has the background assumption that unchallenged assertions are added to the collectively acceptable propositions from which inferences and deliberations are endorsed. See, for a defense of this thesis (Goldberg, 2018)

awful, and everything they say just sounds dumb. Such a person could say to you that *the sky is blue*, but you'll say that this is a hasty generalization from clear days: sometimes it's overcast, so the sky is *grey*, and sometimes the sun is setting, and then it's *red* or *purple*, and sometimes it's night and then the sky is *black*. You see how it goes: you just don't like this guy, so you'll rack your brain to think of ways to say he's wrong, because *you want him to be wrong*. Folks like that just aren't allowed to be right, and for sure, we are going to find ways to keep him from claiming it. What *ad hominem* arguments amount to, then, is either establishing or reminding one's audience that one's opposition belongs to that class of folks who are contemptible, so can't be allowed to be right.⁵ In fact, given that in the case of third-person *ad hominem*, the maligned party isn't even present; it's trash talk without the targets there to defend themselves. And notice that, especially when the audience already agrees with the critical line to take and shares the personal contempt for the opposing arguer (and their type), the case is particularly effective. Further, there is a second piece of interpersonal communication in those third personal cases, since the onlooking audience not only shares in the contempt for the opposition, but it is shown how they each will be treated (and will deserve to be treated) were they to express sympathy for views of the opposition. So, not only is third-person *ad hominem* a strategy of defaming a dialectical opponent so that they will not receive a fair hearing, but it promises similar treatment for those who deviate from the dominant message. In short, *ad hominem* not only demeans one's opposition but it serves as a means for policing one's allies.

A further edge to the third-personal *ad hominem* is worth highlighting, if I am right that it is a negative profile in authority. With wider cultural rifts, the invocation of some group-salient abuse is a restatement of a broader disagreement. Take, for example, the divide between progressives interested in social justice and cultural conservatives who find the proposals of change not only ill-conceived but laughable. The conservative might call the audience's attention to the progressive's "blue hair and pronouns." The progressive may call the conservative a "Christofacist" or "uneducated bumpkin." In these cases, the abuse is not merely some purely irrelevant feature of the speaker, but a cultural marker of the identities and upstream commitments at stake in the disagreement. To the cultural conservative, "blue hair and pronouns" invokes the idea that progressive politics has no appreciation for natural things and seeks to take what is unusual and make it normative. And so, the cultural contempt on the issues is made salient from the start. And the same goes for "Christofacist" and "bumpkin" which portrays politically religious conservatives as those who seek to make democracy in the shape of their parochial religion, and adds to it the idea that they are not worldly enough to make sense of a complicated global political order. These insults, in this case, are placeholders for larger and more thorny disagreements, and they signal to their preferred audiences how to hear and react to what one's opponent has said (or will say).

⁵ It is in this fact of shared contempt that the pragma-dialectical theory holds that the *ad hominem* consequently breaks the *freedom rule*, that the targets of abuse are not free to express themselves. See (van Eemeren and Grootendorst, 1995a; van Eemeren and Garssen, 2023)

One last observation must be made on the effects of argumentative slander with *ad hominem*. So far, I've focused on the *synchronic* features of the use of the fallacy in a dyadic or polyadic dialogue. There are consequences, inferences, and affective results in the here-and-now. But there are *diachronic* effects, too. One I've already noted in the indirect discourse with one's onlooking audience—attacking an opponent not only communicates contempt for the target, but it also establishes a precedent with the group as to how those who side with them are to be treated. So, *over time*, *ad hominem* arguments have cooling effects on dialectical communities—not only on those who forwarded standpoints yielding abuse, but on the audiences who might be interested in improved versions of those views. And it's worth appreciating how vicious rhetorical burns, when funny or particularly painful, are a kind of *tour de force* of one's mockery arsenal for those others who might get out of line. Not only will the targets of the insults not be enthusiastic about further critical dialogue, but even one's allies will think twice about objections or alternatives.

4.

The *relevance puzzle* is a clear result of the explanatory approach to the effectiveness puzzle. If *ad hominem* arguments are arguments from authority in negative profile, then there should be a relevance symmetry between the two argumentative forms.⁶ In fact, it is a pretty regular phenomenon in textbook treatments of *ad hominem* arguments for authors to pause to note that sometimes the speaker's character is relevant to how we assess their claims.⁷ And in the same way we can unjustifiably give a speaker a credibility excess on the basis of irrelevant character traits (as we see with *ad verecundiam* fallacies), so we may give people credibility deficiencies (as with the *ad hominem* fallacies). Reasoning that since there are instances of deferring to properly-identified authorities, we can also identify when people fail on those criteria.

Arguments from authority proceed from three central assessments of an attestor: *competence*, *sincerity*, and *fairness*.⁸ On the one hand, in order to take a person's say-so as good evidence of there being good evidence, we need to think they are competent on the issue. So, a track record, recognition by other experts, accreditation, and so on, should be recognizable. Further, we proceed on the basis of the thought that our attestors are *sincere*—that they are not hiding something or trying to manipulate us. They believe what they are telling us. Finally, being an unbiased and objective arbiter of the evidence and controversy is very important in an authority. That they have been fair to counter-considerations, giving those who disagree opportunities to say their piece, and have listened, is also important for knowing whether their judgment reflects what their expertise on the issue should select. Evidence that

⁶ The analogy between *ad hominem* and *ad verecundiam* has been noted by (Hamblin, 1970; Brinton, 1985; McMurtry, 1986; Walton, 1998)

⁷ For scholarly observations of the potential for relevance, see: (Hoaglund, 1981; Brinton, 1985; Lagerspetz, 1995; Johnstone, 1996; Wijze, 2003; Metcalf, 2005; Macagno, 2013; Williams and Hample, 2019; Wrisley, 2019)

⁸ This three-part condition is a simplified version of the critical questions to ask of arguments from authority, outlined in (Walton, 1997; Walton, Reed and Macagno, 2008).

someone fails any of these criteria is a defeater for the evidence they provide. So, an expert with views not accepted by other experts, or with a track record of wildly inaccurate and costly predictions, or who took a degree from a fake university, is given lower credence than those on the norm, with good track records, and actual degrees. And the same goes for experts who have a conflict of interest on the matter (say, who stand to gain significantly if we were to accept what they claim) or who seem to change their tune depending on to whom they are speaking—they won't be given the same level of credibility as those who have no conflicts of interest and who have been consistent on the issue.

It is in this symmetry between arguments *ad hominem* and arguments *ad verecundiam* that we see criteria for relevance. Importantly, the relevance is to credibility-grading for the forwarded claims, not to whether the claims are true or false. The justification we might have for accepting the claims is undercut by information about the claimant. So, the first site for clarification with the *ad hominem* is whether we are identifying a reason bearing on a claim's truth value based on the character flaw of the claimant or a reason bearing on what degree of support the claiming (given who the claimant is) gives to it. Knowing that someone has a conflict of interest on an issue does not make what they claim *false*, but it does mean that, ideally, one should seek other sources of information to corroborate before one accepts what they have said (Walton, Reed and Macagno, 2008; Jason, 2011).

This relevance, credibility, and fairness approach makes good sense on the symmetry between *ad verecundiam* and *ad hominem*, but there is one important dissimilarity. Arguments from authority are instances from simple attestation, and though many *ad hominem* arguments bear on testimonial contributions, they are also deployed against *arguments*. In testimonial cases, the say-so is the only supporting reason, and so we can then see why assessing speaker character is a plausible path for thinking critically about acceptability. But with arguments, we have further reasons beyond the mere attesting given. So attacks on testimonial sincerity or competence do not touch argumentative quality. So *ad hominem* challenges to, for example, a valid syllogism do not undercut its validity (Jason, 2011). Nor would appealing to the speaker's expertise improve its validity. The whole point of giving an argument is that interlocutors do not have to take our say-so as the primary reason, but we have shared reasons to consider. Sharing the reasons, in this process, makes them *our reasons*, and so they go a long way (ideally) to resolving the issue and taking us to the truth independently of the individuals in the exchange. Arguments, in the end, are designed to be speaker-neutral, so *ad hominem* considerations, once the arguments are out for evaluation, are clearly irrelevant (see Howell and Kingsbury, 2013).

But character can still be relevant, even with arguments. One face of that relevance is in how we might assess validity. In a complicated path through a series of implications, a reasonable question can be how likely the argument is a mere masquerade of validity. The "logic" of someone's argument presented as though it's very tight reasoning can be a rhetorical flourish itself, and if we know that our interlocutor may be insincere in walking us through the inferences, we have reason to be unsure of how tight all this reasoning is. I will admit that there are plenty of times when I have read some work from folks where I was more sure that the technical

vocabulary they had used to explain the validity of their reasoning was chosen for its highest likelihood of yielding precisely the results they wanted. All that technical jazz is too often a cover for rationalization.

With non-monotonic inferences, good inferences can be defeated by adding new information. For example, that people on my college campus are young adults carrying book bags is good evidence that they are college students, but if we were to add to our evidence set that they are also wearing ‘visitor’ badges, the inference is no longer so strong. If our interlocutor has given us a non-monotonic argument, we still can ask whether they have suppressed defeating evidence or have exercised proper diligence in seeking out all the relevant information. So, whether they are sincere, competent, and unbiased is still relevant, even when they’ve given an argument.

Finally, the issues deserve the right degrees of scrutiny, and someone whose values are misaligned may not be the right person for us to trust to be reasoning reliably about it. So, someone who does not suffer if a policy is not effective and stands to gain if it is ratified, regardless of effectiveness, may need to be double-checked on their case for the policy even if there is argument and data in its favor. That their idea looks plausible may be enough for them, but what about people whose lives hang on the results?

The picture emerging here is that, under a wide variety of conditions, the quality of the arguer is still a relevant consideration even when we are assessing the arguments, because we naturally and appropriately ask whether the argument given genuinely reflects the available evidence or is manipulated to garner our assent. So, questions of the virtue or vice of our interlocutor are reasonable when we must trust their representations of the standing evidence, what the dialectical considerations are, whether some errors are more costly than others, and who stands to benefit from the agreements or continued dissensus. It’s obvious that a person’s room-clearing halitosis is not a good reason to think this of their arguments, but concerns about their honesty, good will, fairness, or negligence can be.⁹ It’s worth noting the limits of this point about relevance. The point that who the arguer is does not affect the validity or soundness of the argument still stands, so the speaker-neutrality point is correct. However, how we have assessed the overall quality is speaker-relevant, in particular when we have questions about the speaker’s argumentative character. In cases where speaker sincerity, competence, or fairness is in question, we have reasons undercutting our assessments of argument quality when we must rely on the speaker for the assessments. *Ad hominem* arguments don’t show that an interlocutor’s views are false, but they can show that initial appearances of argument quality can be misleading if not in error.

My point here is that *ad hominem* considerations can be relevant, even when arguments are being considered, because who gives the argument can be a determining factor in how well we’ve assessed whether the argument’s premises

⁹ See (Aberdein and Cohen, 2024) for an overview of virtue theories of argument, making the case that arguer-quality is a relevant consideration for assessing argument-quality. Earlier, Aberdein (2014) had made the case that *ad hominem* objections trade on an ambiguity with relevance of character, showing that who is speaking, under the right conditions, can be a matter of relevance.

support the conclusion and whether we think the premises themselves plausibly represent the available evidence and account for the stakes in the case.

5.

The meta-argumentative puzzle with the *ad hominem* has a number of levels. This shouldn't come as a surprise, since once we go meta with most normative enterprises, the levels pile up quickly. Witness the seemingly endless leveling up in metaphilosophy or the self-referential levels of meta-theater and meta-fiction. To start, meta-arguments are arguments about arguments, given for the sake of evaluating, explaining, or articulating their significant features (Aikin and Casey, 2022a). Notice that my opening cases of popular figures and politicians remarking about the *ad hominem* fallacy was posited not just on the observation that there are widespread instances of the argument form, but that observations of and complaints about the fallacy are widespread, too. We not only reason about things, we reason about each other, and we also reason about our reasoning about things and each other, too. That's the business.

The first edge of the meta-argumentative puzzle is the irony that the *ad hominem* is both widespread in use and in complaint. It's almost as though we know that it's fallacious, but we still do it, and then complain about it. Take the case of former President Barack Obama. He famously complains about the fact of *ad hominem* arguments being so widespread in American political discussion, but his opponents also think he is regularly guilty of the fallacy. Peter Wehner at *Commentary* magazine even went on to title him "*The Ultimate Ad Hominem President*" (2013). It's equally true that Ann Coulter, whose book *Slander* is posited on the thought that liberals use the *ad hominem* as their default argumentative move, tends to use the tactic as a favorite herself. Take the simple fact that one of her follow-up books, *Godless*, is an argument that liberal sensibilities are founded on sinful opposition to the divine. In it, she identifies liberals as "ugly feminists," and "pompous idiots," and she says that after she criticized Hillary Clinton, "Hillary beat a hasty retreat on her chubby little legs" (2007, 289). For sure, all the complaints in the literature broadly are in the form of identifying *ad hominem* in the *third person* form, as opposed to the *first personal*—namely, that *they* use the *ad hominem*, not that *I* or *we* use it. So there must be another distinction, but this time at the meta-level. In this case, there is the difference, on the meta-argumentative level of assessing arguments as arguments, between:

First-person: I am giving a fallacious *ad hominem* argument, and

Second-person: You are giving a fallacious *ad hominem* argument, and

Third-person: They are giving a fallacious *ad hominem* argument.

That is, it seems conflicted for someone to say, "*I* sure do give a lot of fallacious *ad hominem* arguments," but it certainly has no tension for someone to say, "*They* use a lot of *ad hominem* reasoning," or "That argument *you* gave was an *ad hominem* fallacy." The lesson, of course, is that fallacious inferences generally are not

identifiable in the first-person: what it is to make an inference is to endorse the support the movement purports to trace. So, seeing an argument as fallacious and endorsing it at the same time has the distinct flavor of an inferential version of Moore's paradox (where one says that one believes a proposition but then says that the proposition isn't true). Understanding fallacies, then, usually requires a double-vision in appreciating what is tempting about the inference but also in seeing it as a bad inference—seeing one facet clearly usually occludes and precludes the other. The result is that one is bound to see the fallacies committed by one's opponents, but when one's own arguments take *ad hominem* form, one will see those as legitimate critical questions about speakers and their arguments. So, the fallacious irrelevant cases will all feel, in the first person, relevant.

The second level to the meta-argumentative puzzle is that the observations and complaints about *ad hominem* fallacies (in the third person) prompt a series of new meta-arguments. These are initiated by the thought that if information about a speaker's character can be meta-evidence about the quality of their argumentative performance, then fallacious arguments given by the opposition can be evidence about the quality of the evidence at their disposal. Again, meta-arguments, as interpersonal inferences, require some counterfactual reasoning as we saw with other forms of interpersonal inference with pragmatic implicature. But in these cases, we make inferences about the quality of the opposition's overall case from the fact that they argue in the form of the *ad hominem* attack. The meta-argument comes in a modest and a strong form.

The modest form of the meta-argument is a *defeating reasons interpersonal inference*. The core thought behind this inference is that if name-calling and insults are the best the opposition has, then that is evidence about what kind of case they have. It's evidence that their case is weak. Here the argument is in rough form:

The opposition uses fallacious *ad hominem* arguments

Were they to have better reasons, they would have given them. But they didn't.

So, the *ad hominem* attack is the best case they have.

The important assumption here is that we, if we are following the norms of communication, have good reason to give our best reasons, and given the stakes, it's incumbent for us to work hard for those reasons to be manifestly good. So, the meta-argument, in its defeating reasons form, runs from some counterfactual reasoning about one's opposition and from the fact that they give a lot of otherwise bad arguments. But there is another move for the meta-arguer. The next step is to see that the defeating version of the meta-argument can, by a second contrast, provide further reason to hold that one's own reasons are quite good. Were our own case worse off, the other side (given the communicative obligations and the stakes for the discussion) would have much to say in criticism of our first-order arguments. So, if their primary criticisms are about our bad breath or otherwise

to insult us, we have reason to see that our own case is quite good. So, the stronger version of the other-regarding meta-argument has an amplifying effect:

The opposition uses fallacious *ad hominem* arguments

Were they to have better critical replies to our case, they would use them.
But they didn't.

So, the *ad hominem* attacks are indirect confirmation that we have a good case.

The reality is that the *ad hominem* is old news to people who think about critical thinking. The issue is that we still use personal attack despite knowing it to be fallacious, and so we proceed to give our opponents reason to make inferences about our own cases and their own in the process. We get the meta-argumentative puzzle because we think our own *ad hominem*s are relevant, so we use them. Or we give them behind what we think are closed doors just as expressions of shared contempt, but our slanders get back to our targets. And *they* don't see them as relevant or just behind-closed doors joshing. Rather, because they see the *ad hominem* attacks on them as not only fallacious but telling, this fact is relevant as meta-argumentative evidence for them. The problem is not that there isn't enough critical thinking, but that critical thinking runs amok.

Consider this instance. When the question, "Does using *ad hominem* make your argument weak?" was asked on *Quora*, an aggregator blog, commenter Chris Tor (identifying himself as a "person of the world") replied:

Resorting to *ad hominem* means you have given up on making an argument. This is usually because you don't have a good one, but it can also just come from poor argumentation skills. (*Does resorting to ad hominem mean your argument is weak?* - *Quora*, no date)

This line of reasoning has been taken on by all sides in political debates. Consider the following from the American Right:

Ben Shapiro: "These bullies – they're not going to use things like logic to argue their case. They're simply going to slander people." (Shapiro, 2013)

Kirsten Powers: "It's part of a larger effort to demonize ... anyone who doesn't agree.... Their goal is to shut down debate they fear they are losing on the merits." (Powers, 2015)

Ann Coulter: "Progress cannot be made on serious issues because one side is making arguments and the other side is throwing eggs.... Logic is not their *métier*. Blind religious fervor and denigration is." (Coulter, 2002)

And here is one from former President Obama, commenting on the poor quality of debate among the 2016 Presidential hopefuls:

[Y]ou've heard me ... try to get a good argument on the other side that's based in fact as opposed to rhetoric. And I haven't gotten one yet. So if you're asking me, how do you think our argument is going, it's going great.... We have robust debates, we look at the facts, there are going to be disagreements. But we just don't fling out ad hominem attacks like that, because it doesn't help inform the American people.(2015)

In a way, the meta-argumentative inferences, in both modest and strong forms, are plausible explanations for why argumentative trolling works the way it does. One annoys one's opponent to the point where they lose their temper and call you a name. Then you use this as evidence of their intellectual vice, failure of reason, and the comparative quality of your preferred opposing view.

The third level to the meta-argumentative puzzle is about the meta-meta-evidential problem we face if we are tempted by the defeating or amplifying meta-arguments. Moreover, this consideration is a meta-evidential defeater even for those who complain about first-order *ad hominem* arguments. The defeating consideration is that of what's called the 'nasty effect' that bears on those who have witnessed abusive language in the midst of a critical discussion: they end up more polarized on the issues but recall fewer details of the matter beyond the insults. In a study by Anderson *et al.* (Anderson *et al.*, 2014), readers were given an article with balanced content on the pros and cons of nanotechnology, with one group also reading hostile comments on the piece, and with another group not exposed to the comments. The result was that those who had prior support for one view or another were polarized by the comments, compared to the group not exposed to the comments. In short, those who had a view before reading the nasty comments about those who hold that view were made more extreme in their views. The lesson is that nasty comments polarize audiences who'd already chosen sides, making them more ardent supporters of the views with which they'd arrived. It's not that they became more informed on the issue, but more exercised by the incivility of opposition to their views.

The meta-meta-argumentatively defeating problem is that if we, as partisans, complain about the *ad hominem* arguments given by the other side on the first-order issue, and we are tempted by the meta-argumentative inferences on the second level, *then it's likely that we've also been exposed to the conditions constituting the nasty effect.* It's likely that our own confidence in our own side in the debate is a product of the polarizing forces of incivility in the exchange as it is of the force of evidence. It's true also for those on the other side, too, but the point here is that *our own* concerns and complaints about the *ad hominem* attack is evidence that we've been exposed to the polarizing conditions. (In a way, the complaint about the other side, in second- and third-personal meta-argumentative form, making use of and making inferences about *ad hominem* has its own valence of this concern.)

A further problematizing concern is that given belief-bias (the phenomenon wherein believing the conclusion of an argument makes it more likely one will think the reasoning is good), even discriminating the irrelevant cases of our own

uses of *ad hominem* reasoning will be undercut. Given this bias, the opponent's *ad hominem* will overwhelmingly appear fallacious, and *ad hominem* arguments from one's own side will appear relevant.¹⁰ A sports analogy might be useful here. When one's team is engaged in a particularly physical game with a heated rival, there are often many fouls. But partisans tend to see the opposition committing more fouls than their own team (Hastorf and Cantril, 1954). Selective perception and memory, driven by my-side bias, then stands as a meta-meta-evidential defeater for these tempting inferences. Take the game, on the analogy, to be arguments, the fouls to be fallacies, and team-loyalty to be prior commitment, and we can see the problem in high definition. In short, if you're inclined to perform one of the *ad hominem* meta-arguments, you have reason to distrust your assessment of the meta-evidence, because you've likely been exposed to the conditions constituting the nasty effect.

This third level of the meta-argumentative puzzle, then, works as a kind of epistemic and argumentative tragedy—as we have become more sensitive to abuses of argument and capable of reflection upon it and those who argue, we have deepened our capacity to rationalize our commitments and polarize ourselves.¹¹

6.

A lesson of this long arc through the *ad hominem* is that norms of argumentative civility are what they are for a reason. They are not mere impositions of etiquette or cultural gatekeeping, but rather, they are means of staying on task and preventing us from being distracted in the midst of a discussion. Argumentative exchange is an interpersonal interaction rife with triggers for escalation, and the *ad hominem*, even when it is relevant, supercharges this effect.¹² Given the risks of irrelevance on the first order, the meta-argumentative inferences that interlocutors are tempted to make on the second order, and the meta-meta-argumentative defeat on the third level of *ad hominem* arguments, it seems to me the best policy is to avoid giving these kinds of arguments if one has other evidential paths. And it would behoove us to disincentivize others from using them, too. This is, as John Woods dismissively terms it in his case for easing up on the *ad hominem*, a “goody-two-shoes” approach to argument, but given our aspirations of resolving disagreements and getting to the truth, this seems the most prudent path. Yes, *ad hominem* arguments can be relevant, but the odds are that you'll over-attribute relevance when the arguments undercut people with whom you disagree and under-attribute relevance when they are deployed against views and people you like. This phenomenon of *ad hominem*

¹⁰ Assuming that belief-bias is a particular version of my-side bias, the simple version is that it's easier to see fallacies in arguments for views one rejects than in those one supports. This seems even more pronounced when the contributions can be interpreted as *attacks*. The place to start with this interpretive element to attributing *ad hominem* is that identifying instances in opposition is often a function of affective polarization. See (Heinzmann and Tran, 2024).

¹¹ The “Owl of Minerva Problem” is one that dogs attempts by those wielding meta-level discourse hoping to improve first-order discourse. See (Aikin, 2020; Godden, 2022; Aberdeen, 2023)

¹² For an example of a “cascade of hate,” see what Álvarez-Benjumea observes when hateful content is given voice in deliberative spaces (Álvarez-Benjumea, 2023)

abuse polarizes us in a way that makes us less likely to reach agreement or follow the evidence. The result is that if only “jerks” use the *ad hominem* (Verber 2012), it may take a real asshole, after knowing the risks, to endorse it as appropriate argumentation.

SCOTT AIKIN
VANDERBILT UNIVERSITY
scott.f.aikin@vanderbilt.edu

Acknowledgements. Thanks go to John Casey, David Godden, Tempest Henning, Robert Talisse, Lucy Alsip Vollbrecht, and the audience at Michigan State University’s Asquith Symposium on Informal Logic for feedback on an early draft of this essay. Special thanks go to this journal’s two anonymous reviewers for very helpful recommendations for improvement.

References

- Aakhus, M. and Lewiński, M. (2023) *Argumentation in Complex Communication*. New York: Oxford University Press.
- Aberdein, A. (2014) ‘In Defence of Virtue: The Legitimacy of Agent-Based Argument Appraisal’, *Informal Logic*, 34(1), pp. 77–93. Available at: <https://doi.org/10.22329/il.v34i1.3938>.
- Aberdein, A. (2023) ‘The Fallacy Fallacy: From the Owl of Minerva to the Lark of Arete’, *Argumentation*, 37(2), pp. 269–280. Available at: <https://doi.org/10.1007/s10503-023-09595-9>.
- Aberdein, A. and Cohen, D.H. (2024) ‘Virtue Theories of Argument’, *Inquiry: Critical Thinking Across the Disciplines*, 33(2), pp. 117–142. Available at: <https://doi.org/10.5840/inquiryct20244115>.
- Aikin, S. (2020) ‘The Owl of Minerva Problem’, *Southwest Philosophy Review*, 36(1), pp. 13–22. Available at: <https://doi.org/10.5840/swphilreview20203613>.
- Aikin, S. and Casey, J. (2022a) ‘Fallacies of Meta-argumentation’, *Philosophy and Rhetoric*, 55(4), pp. 360–385. Available at: <https://doi.org/10.5325/philtrhet.55.4.0360>.
- Aikin, S. and Casey, J. (2022b) *Straw Man Arguments: A Study in Fallacy Theory*. London: Bloomsbury.
- Aikin, S. and Talisse, R. (2019) *Why We Argue (and How We Should): Guide to Disagreement in an Age of Unreason*. 2nd edn. New York: Routledge.
- Álvarez-Benjumea, A. (2023) ‘Uncovering hidden opinions: social norms and the expression of xenophobic attitudes’, *European Sociological Review*, 39(3), pp. 449–463. Available at: <https://doi.org/10.1093/esr/jcac056>.
- Anderson, A.A. et al. (2014) ‘The “Nasty Effect:” Online Incivility and Risk Perceptions of Emerging Technologies’, *Journal of Computer-Mediated Communication*, 19(3), pp. 373–387. Available at: <https://doi.org/10.1111/jcc4.12009>.
- Bowell, T. and Kingsbury, J. (2013) ‘Virtue and Argument: Taking Character Into Account’, *Informal Logic*, 33(1), pp. 22–32. Available at: <https://doi.org/10.22329/il.v33i1.3608>.
- Brinton, A. (1985) ‘A rhetorical view of the ad hominem’, *Australasian Journal of Philosophy*, 63(1), pp. 50–63. Available at: <https://doi.org/10.1080/00048408512341681>.
- Brinton, A. (1995) ‘The Ad Hominem’, in H.V. Hansen and R.C. Pinto (eds) *Fallacies: Classical and Contemporary Readings*. University Park, Pa: Pennsylvania State University Press, pp. 213–221.
- Budzynska, K. and Witek, M. (2014) ‘Non-Inferential Aspects of Ad Hominem and Ad Baculum’, *Argumentation*, 28(3), pp. 301–315. Available at: <https://doi.org/10.1007/s10503-014-9322-6>.
- Casey, J. (2022) ‘Beliefs, Commitments, and Ad Baculum Arguments’, *Languages*, 7(2), p. 107. Available at: <https://doi.org/10.3390/languages7020107>.
- Coulter, A.H. (2002) *Slander: liberal lies about the American right*. 1st pbk. ed. New York: Three Rivers Press.
- Coulter, A.H. (2007) *Godless: The Church of Liberalism*. New York: Three Rivers Press.

- Does resorting to ad hominem mean your argument is weak? - Quora (no date). Available at: <https://www.quora.com/Does-resorting-to-ad-hominem-mean-your-argument-is-weak> (Accessed: 30 April 2024).
- van Eemeren, F.H. and Garssen, B. (2023) 'The Pragma-Dialectical Approach to the Fallacies Revisited', *Argumentation*, 37(2), pp. 167–180. Available at: <https://doi.org/10.1007/s10503-023-09605-w>.
- van Eemeren, F.H. and Grootendorst, R. (1995a) 'Argumentum Ad Hominem: A Pragma-Dialectical Case in Point', in H.V. Hansen and R.C. Pinto (eds) *Fallacies: Classical and Contemporary Readings*. University Park, Pa: Pennsylvania State University Press, pp. 223–228.
- van Eemeren, F.H. and Grootendorst, R. (1995b) 'The Pragma-Dialectical Approach to Fallacies', in H.V. Hansen and R.C. Pinto (eds) *Fallacies: classical and contemporary readings*. University Park, Pa: Pennsylvania State University Press, pp. 130–144.
- Godden, D. (2022) 'Getting Out in Front of the Owl of Minerva Problem', *Argumentation*, 36(1), pp. 35–60. Available at: <https://doi.org/10.1007/s10503-021-09554-2>.
- Goldberg, S. (2018) 'Dissent: Ethics and Epistemology', in C.R. Johnson (ed.) *Voicing Dissent*. New York: Routledge, pp. 40–60.
- Hamblin, C.L. (1970) *Fallacies*. London: Methuen.
- Hastorf, A.H. and Cantril, H. (1954) 'They saw a game; a case study', *The Journal of Abnormal and Social Psychology*, 49(1), pp. 129–134. Available at: <https://doi.org/10.1037/h0057880>.
- Heinzelmann, N. and Tran, V. (2024) 'Extremists are more confident', *Erkenntnis*, 89(5), pp. 2031–2056. Available at: <https://doi.org/10.1007/s10670-022-00616-9>.
- Hinman, L.M. (1982) 'The Case for Ad Hominem Arguments', *Australasian Journal of Philosophy*, 60(4), pp. 338–345. Available at: <https://doi.org/10.1080/00048408212340741>.
- Hitchcock, D. (2017) *On reasoning and argument: Essays in informal logic and critical thinking*. Cham, SU: Springer. Available at: <https://doi.org/10.1007/978-3-319-53562-3>.
- Hoaglund, J. (1981) 'Argumentum Ad Hominem: Aut Bonum Aut Malum?', *Informal Logic*, 4(3). Available at: <https://doi.org/10.22329/il.v4i3.2773>.
- Hundleby, C. (2023) 'Social Justice, Fallacies of Argument, and Persistent Bias', *Argumentation*, 37(2), pp. 281–293. Available at: <https://doi.org/10.1007/s10503-023-09603-y>.
- Jason, G.J. (2011) 'Does Virtue Epistemology Provide a Better Account of the Ad Hominem Argument? A Reply to Christopher Johnson', *Philosophy*, 86(1), pp. 95–119. Available at: <https://doi.org/10.1017/S0031819110000616>.
- Johnstone, H.W. (1996) 'Locke and Whately on the Argumentum ad Hominem', *Argumentation*, 10(1), pp. 89–97. Available at: <https://doi.org/10.1007/BF00126161>.
- Lagerspetz, E. (1995) 'Ad Hominem Arguments in Practical Argumentation', *Argumentation*, 9(2), pp. 363–370. Available at: <https://doi.org/10.1007/bf00721966>.
- Lewiński, M. (2019) 'Argumentative Discussion: The Rationality of What?', *Topoi*, 38(4), pp. 645–658. Available at: <https://doi.org/10.1007/s11245-015-9361-0>.
- Lewiński, M. and Aakhus, M. (2014) 'Argumentative Polylogues in a Dialectical Framework: A Methodological Inquiry', *Argumentation*, 28(2), pp. 161–185. Available at: <https://doi.org/10.1007/s10503-013-9307-x>.
- Macagno, F. (2013) 'Strategies of Character Attack', *Argumentation*, 27(4), pp. 369–401. Available at: <https://doi.org/10.1007/s10503-013-9291-1>.
- McMurtry, J. (1986) 'The Argumentum Ad Adversarium', *Informal Logic*, 8(1), pp. 27–36.
- Metcalf, R. (2005) 'Rethinking the Ad Hominem: A Case Study of Chomsky', *Argumentation*, 19(1), pp. 29–52. Available at: <https://doi.org/10.1007/s10503-004-2069-8>.
- 'Obama is the Ultimate Ad Hominem President' (2013) *Commentary Magazine*, 15 May. Available at: <https://www.commentary.org/peter-wehner/obama-is-the-ultimate-ad-hominem-president/> (Accessed: 7 May 2024).
- Powers, K. (2015) *The silencing: how the left is killing free speech*. Washington, D.C: Regnery Publishing.
- Remarks by President Obama and Prime Minister Hailemariam Desalegn of Ethiopia in Joint Press Conference (2015) [whitehouse.gov](https://obamawhitehouse.archives.gov/the-press-office/2015/07/27/remarks-president-obama-and-prime-minister-hailemariam-desalegn-ethiopia). Available at: <https://obamawhitehouse.archives.gov/the-press-office/2015/07/27/remarks-president-obama-and-prime-minister-hailemariam-desalegn-ethiopia> (Accessed: 30 April 2024).

- Rubin, J. (2021) 'CPAC missteps again [UPDATED]', *Washington Post*, 2 December. Available at: <https://www.washingtonpost.com/blogs/right-turn/wp/2021/03/05/cpac-missteps-again/> (Accessed: 6 May 2024).
- Sen. Cruz Slams Dr. Fauci on 'Hannity': 'The Most Dangerous Bureaucrat in the History of the Country' | U.S. Senator Ted Cruz of Texas (2021). Available at: <https://www.cruz.senate.gov/newsroom/press-releases/sen-cruz-slams-dr-fauci-on-hannity-the-most-dangerous-bureaucrat-in-the-history-of-the-country> (Accessed: 6 May 2024).
- Shapiro, B. (2013) *Bullies: how the left's culture of fear and intimidation silences America*. New York: Threshold Editions.
- Walton, D.N. (1997) *Appeal to Expert Opinion*. University Park, Pa: Pennsylvania State University Press.
- Walton, D.N. (1998) *Ad hominem arguments*. Tuscaloosa: University of Alabama Press.
- Walton, D.N., Reed, C. and Macagno, F. (2008) *Argumentation schemes*. Cambridge; New York: Cambridge University Press.
- Wijze, S. de (2003) 'Complexity, Relevance and Character: Problems with Teaching the Ad Hominem Fallacy', *Educational Philosophy and Theory*, 35(1), pp. 31–56. Available at: <https://doi.org/10.1111/1469-5812.00004>.
- Williams, D.C. and Hamble, D. (2019) 'Rising to the defense of ad hominem arguments', in C. Winker (ed.) *Networking Argument*. New York: Taylor & Francis, pp. 427–432.
- Woods, J. (2007) 'Lightening up on the Ad Hominem', *Informal Logic*, 27(1), pp. 109–134. Available at: <https://doi.org/10.22329/il.v27i1.467>.
- Wrisley, G. (2019) 'Ad Hominem: Direct', in R. Arp, S. Barbone, and M. Bruce (eds) *Bad Arguments*. Hoboken (N.J.): Wiley Blackwell, pp. 83–6.
- Yap, A. (2013) 'Ad Hominem Fallacies, Bias, and Testimony', *Argumentation*, 27(1), pp. 97–109.
- Yap, A. (2015) 'Ad Hominem Fallacies and Epistemic Credibility', in C. Dahlman and T. Bustamante (eds) *Argument Types and Fallacies in Legal Argumentation*. Springer Verlag.