

A Modal Criterion for Epistemic Argumentation

JOB DE GREFTE

*Department of Philosophy
University of Groningen
9712 CP Groningen
Netherlands
j.a.m.de.grefte@rug.nl*

Abstract: In this paper, I adopt a pluralistic framework on argumentation, where the norms governing argumentation depend on the aim with which we engage in the practice. A domain of specifically epistemic argumentation is singled out, and I argue based on recent findings in modal epistemology that this domain is governed by the modal norm of safety; where a belief is safe just in case it is produced by a method that would not easily produce a false belief. While this criterion is well-known and uncontroversial in epistemology, it has hitherto not been applied to epistemic theories of argumentation. I show the fruitfulness of bringing this modal norm into our theory of argumentation by arguing that this allows for a novel and superior perspective of the relevance of the persistent interlocutor in argumentation theory, and on the relation between dialectical and epistemic norms more generally.

Résumé: Dans cet article, j'adopte un cadre pluraliste sur l'argumentation, où les normes qui dirigent la construction et l'évaluation de l'argumentation dépendent du but de notre engagement dans cette pratique. Un domaine d'argumentation spécifiquement épistémique est distingué, et je soutiens, sur la base de découvertes récentes en épistémologie modale, que ce domaine est dirigé par la norme modale de sécurité, selon laquelle une croyance est sûre juste au cas où elle serait produite par une méthode qui ne produirait pas facilement une fausse croyance. Bien que ce critère soit bien connu et non controversé en épistémologie, il n'a jusqu'à présent pas été appliqué aux théories épistémiques de l'argumentation. Je montre la fécondité d'introduire cette norme modale dans notre théorie de l'argumentation en soutenant que cela permet une perspective nouvelle et supérieure sur la pertinence de l'interlocuteur persistant dans la théorie de l'argumentation, et plus généralement sur la relation entre les norms dialectiques et épistémiques.

Keywords: epistemic argumentation, safety, persistent interlocutor

1. Introduction

We argue with each other all the time, but what is good argumentation? Rhetorical (Perelman and Olbrechts-Tyteca 1969), pragma-dialectical (Eemeren and Grootendorst 2003) and epistemological (Goldman 2003; Lumer 2005) approaches each identify different norms for good argumentation. Usually, these approaches are seen as competitors (e.g. Lumer 2005). I have two aims in this paper. First, to paint a different picture; an inclusive approach to argumentation that conceptualizes argumentation as a multi-purpose tool subject to different norms depending on contextually determined aims. On this view, rhetorical, pragma-dialectical and epistemological approaches peacefully coexist.

My second aim in this paper is to apply the inclusive approach by delineating a specific form of argumentation for which the epistemological approach is most suitable, and to contribute to the existing literature on epistemological approaches by incorporating insights from recent (modal) epistemology. Along the way, we'll see how this approach may provide novel insights on a notorious character in argumentation theory: the persistent interlocutor.

The paper is structured as follows. In Section 2, I will argue for my inclusive approach to argumentation theory, and carve out a specifically epistemic form of argumentation. In Section 3, I discuss the safety-based account of knowledge. In Section 4, I apply the safety condition by developing a safety-based condition for argumentation, and show how argumentation may increase the safety of our beliefs. In Section 5, I show how the approach provides novel insight on the persistent interlocutor. In Section 6, I consider objections, and Section 7 concludes.

2. An inclusive approach to argumentation

According to rhetorical approaches, argumentation serves to persuade. According to pragma-dialecticians, the aim of argumentation is to resolve differences of opinion. According to epistemological theories of argumentation, the aim of argumentation is to produce justified belief or knowledge. These approaches are usually seen as competitors in the sense that only one of them can be right (Lumer 2005).

Here, I will resist this picture. In particular, I will draw distinctions between arguments and argumentation, between descriptive and normative approaches to argumentation and explain my own inclusive approach to argumentation.

The first distinction to draw is between arguments and argumentation (Dutilh Novaes 2021), or as the distinction is sometimes labelled, arguments as *products* vs arguments as *processes* or argument₁ vs argument₂ (O’Keefe 1992). Here, I will adopt the former terminology. Arguments, as I will conceive them, are sets of statements, such that some of them support others. I am explicitly not opting for a characterization of arguments in terms of (sets of) propositions, because I think arguments require intentionality; propositions that have never been *used* to support each other are not yet arguments.

Argumentation, then, is the process of producing such arguments. In this paper, I am interested primarily in social settings: while it may be possible to “argue with oneself,” I will focus on the norms governing settings where we argue with each other. I remain uncommitted as to the question whether these norms generalize to intrapersonal settings.¹

While argumentation can of course occur in much larger groups, the central case in this paper will be a setting with two agents arguing with each other. Let us call these agents agent P and agent O (for proponent and opponent). P and O could take various roles, but in its simplest form, P presents a thesis and produces an argument for it in an effort to convince O. The role of O is to ask critical questions regarding P’s argument—does the conclusion really follow; are the premises plausible?—this kind of argumentation is known as a simple, or non-mixed dispute (van Eemeren and Grootendorst 2016, Sect. 2.3). If O defends theses of her own, she too engages in argumentation and the dispute becomes mixed. For reasons of simplicity, we will focus on non-mixed disputes in this paper.

Argumentation thus, is a practice. This practice is goal-directed: people use argumentation for a reason. To be sure, this reason can

¹ There is some reason to suspect they might, however. As Catarina Dutilh-Novaes has argued, argumentative practices may influence our capacity for deductive reasoning (Dutilh Novaes 2013).

be as simple as to have some fun or spend some time, but argumentation takes some effort so people do not engage in it without some expected benefit (although they may be drawn into argumentation by others for *their* benefit). Because argumentation is a goal-directed practice, we may study it descriptively or normatively. That is, we may try to describe how people actually argue, or in conjunction, we may try to investigate how people can achieve their argumentative aims effectively—and provide norms for *good* and *bad* argumentation.

This paper clearly falls on the normative side of this distinction. I will posit some (hopefully uncontroversial) aims with which people engage in argumentation, and investigate what it would take to achieve these aims effectively, thus deriving norms for good argumentation.

The pluralism of my account follows from the fact that people may use argumentation for different aims, and that different aims will produce different norms. For example, argumentation may be used by politicians to persuade people that a certain course of action should be taken. In such a setting, effectively achieving this goal may require great rhetorical skill, and even an occasional liberal interpretation of the truth. On the other hand, argumentation in science is often used to discover unknown truths, and with such an aim a strict truth norm may be more effective. What this means is that we must relativize our norms of argumentation to the aims with which we engage in the practice, and that different norms will be associated with different aims. Argumentation is a tool, and whether a specific use of that tool can be criticized depends on what we aim to do with it.

In this paper, the aim that I want to focus on is epistemic, namely the acquirement of knowledge. This aim is not the same as the aim of finding truth, since as we will see below, knowledge requires more than just truth. But ambitious as the aim may be, I submit it is one of the more important aims of argumentation. In science, and in daily life as well, we often aim at more than just luckily true belief: we aim at knowledge. And we use argumentation as a tool to acquire it. Let us call argumentation aimed at knowledge epistemic argumentation. In the remainder of this paper, I will investigate one of the main norms governing epistem-

ic argumentation. To that end, we will have to look in the next section at the requirements for knowledge.

3. The safety-based account of knowledge

What is clear is that at least sometimes, we use argumentation to provide us with knowledge. The primary example here is science, where debates in scientific journals and conferences can be seen as temporally extended processes of argumentation with a clear epistemic aim: scientists debate hypotheses at least partly, we may assume, in order to know whether they are true. I have called this kind of argumentation *epistemic argumentation*. An account of epistemic argumentation requires an account of what it is such argumentation tries to achieve: knowledge. To be sure, the vast literature on this subject cannot even in main lines be summarized here. I thus opt for a modest strategy: I focus on one recent and popular account to see what would follow regarding the norms on epistemic argumentation. I do not claim this is the only way of conceptualizing knowledge, but the criterion is relatively uncontroversial, and as we will see, it will allow us to paint a plausible and novel picture of the epistemic relevance of the persistent interlocutor.

I focus on what is known as the *safety*-condition on knowledge. We are all familiar with the notion of safety: we can be (relatively) safe from accidents if we properly protect ourselves, or we may be unsafe from such accidents if our behaviour is reckless and we fail to take the proper precautions. To be safe means not to be at risk from some unwelcome event occurring. In the case of knowledge, the focus is on the risk of forming false beliefs. At its most simple formulation, a safety-condition of knowledge says that we know that p is the case (the account is in the first instance an account of propositional knowledge) only if we are not at risk of believing falsely.

A couple of things need to be said to make this rough account more precise. First, what it takes in general to be safe from some risk occurring. Here, safety theorists like Sosa, Williamson and Pritchard take a modal approach, and I will follow their lead in this paper (Pritchard 2005; Sosa 1999; Williamson 2000). A modal account of safety says that we are safe from some risk occurring if that risk does not only fail to occur in the actual world, but also in

most nearby possible worlds, where nearby possible worlds are worlds that are similar to the present one. The more similar the worlds where the risk event occurs are to the actual world, the more easily the risk event can occur, and the less safe we are from it occurring.

Tricky questions arise when we try to make the notion of world similarity more precise: how do we measure overall world-similarity? Do all respects in which worlds can differ from each other matter equally? Here I do not aim to provide an answer to such questions. Rather I take the notion of similarity as primitive; it seems we are able in practice to make fairly clear judgements of similarity about particular cases, and I take these judgements to inform our judgements of safety. For example, we use airbags in our cars because we think, I presume, that with such devices the range of circumstances in which car accidents fail to cause serious injuries (the risk event in question) substantially widens.² On the present analysis of safety, that is why we consider cars with airbags safer than cars without.

The safety-based account thus says that one knows just in case there are no nearby (similar) worlds where one believes falsely. This needs to be qualified, however. Take the following case: I believe falsely that water is H_3O . Now I happen by accident to overhear a chemistry professor explaining that it is, in fact, H_2O . Supposing that not much would need to change to the actual world for me to fail to overhear this professor, there will be similar worlds where I continue to believe falsely that water is H_3O . But still, it seems, by modifying my belief on the basis of reliable testimony from an expert I should be able to acquire knowledge about the chemical nature of water (if this would not be so, many of us would lack this knowledge!). What matters for whether we have knowledge thus is not whether we believe falsely *at all* in nearby worlds, but rather whether we believe falsely on the same basis, using the same method. To illustrate this further: I can come to know about the chemical nature of water on the basis of reliable expert testimony, but not on the basis of reading tea-leaves, even if the tea-leaves happen to predict the nature of water to be H_2O ,

which is true necessarily.³ The reason is that tea-leaf reading would have easily led me to form a different, false belief about the chemical nature of water (presuming, of course, that this method is unreliable and will produce semi-random answers to the questions posed). We thus arrive at the following specification of the safety-based account of knowledge:

K_{safe}: agent S's belief that *p* constitutes knowledge iff the method S used to form her belief that *p* does not produce a false belief in nearby worlds.⁴

One specification is in order before we continue to apply this norm in the next section. For it may be thought that K_{safe} is a reliability condition on knowledge as it relates to the tendency of our methods to produce true belief rather than false. And while I think in a general sense, standard forms of reliabilism and safety-based approaches belong to the same externalist family in epistemology, I will nevertheless here distinguish the two. In this paper, reliability concerns the proportion of cases where a method produces truth, whereas a safety concerns the modal closeness of the nearest world where the method produces error. Lottery cases famously bring out the distinction. If I believe on the tremendously long odds involved that my ticket is a loser, it is reliably formed yet not safe, for only a few different numbers would need to come up for my ticket to be a winner, and so the nearest world where I err is

² This is not the same as saying that safety requires a low probability of the risk event occurring. Take the following example (from Neil Levy): suppose I play Russian roulette with a 100-chamber gun filled with just a single bullet. I am not safe, even if the probability of death is quite small—and we may of course arbitrarily increase the number of chambers. For more on the relation between safety, possible worlds and probability, see de Grefte, J. (2020).

³ Pritchard noted the problems caused by necessarily true beliefs and decided to restrict his account of safety to contingent propositions (Pritchard 2005). Since this move is ad-hoc, a better solution is to allow the formation of false but *different* beliefs in the modal neighbourhood to undermine the safety of one's belief in the actual world, as Pritchard acknowledges in his (2012, pp. 256-57).

⁴ The safety condition is usually understood as a necessary condition on knowledge, rather than as a necessary and sufficient one. For reasons of simplicity, I here favour a necessary and sufficient reading. The points of this paper can also be made on the weaker reading, however.

modally close.⁵ Since lottery cases are prime examples of ignorance rather than knowledge, it seems it is safety rather than reliability that is required for knowledge.

Having explained the safety requirement on knowledge, I develop a safety-based approach to epistemic argumentation in the next section.

4. A safety-based account of epistemic argumentation

In Section 2, I delineated the notion of epistemic argumentation: argumentative processes with the aim of producing knowledge. In the previous section, I explained the safety condition on knowledge. In this section, I will connect these points and develop a safety-based account of argumentation. I contrast the theory to other epistemic theories of argumentation, and explain how argumentation may help us achieve knowledge.

The question before us is how we may gain knowledge through argumentation. If knowledge requires safety, then one way in which argumentation may help us achieve knowledge is by increasing the safety of our beliefs. In this section, I will explain how argumentation may do this under normal conditions. In the next section we will look at the abnormal case of the persistent interlocutor, the epistemic import of which is illuminated with the help of our safety-based account.

Why think argumentation helps increase the safety of our beliefs? We have seen above that safety requires that one's belief-forming process does not easily produce false belief, in the sense that it does not lead one to form false beliefs in nearby possible worlds. To show that argumentation can increase the safety of our beliefs, we thus need to explain how argumentation may increase the sphere of possible worlds in which one believes truly.

The relevant notion of safety applies to belief-forming methods, so the first thing to establish is that argumentation is a belief-forming process at all. A belief-forming process is a process that produces beliefs. If that process occurs within the skull of a single agent, we speak of individual belief-forming processes. If it crucially involves other people, we speak of a social belief-forming

⁵ See (Williamson 2009)

process.⁶ Since the process of argumentation as we have been characterizing it involves different people, it will be a social belief-forming process.

Does argumentation indeed produce beliefs? According to an influential account of argumentation, the aim of argumentation is to resolve a difference of opinion (van Eemeren et al. 2004). Since opinions are plausibly regarded as beliefs here, and since to resolve a difference in belief at least one party to the argument must change her belief, it follows that argumentation is a process that aims to produce new beliefs.

If argumentation is a belief-forming process, how may this process increase the safety of our beliefs? In order to show this, we need to say a bit more on the structure of argumentative exchange. Throughout this paper, I have been drawing on the pragma-dialectical framework of argumentation developed by Frans van Eemeren and Rob Grootendorst (2016, 2004), and I will continue

⁶ As all processes, belief-forming processes can be described on different levels of generality. We may investigate the process of forming beliefs by eyesight, for example, or that of forming beliefs by eyesight in good lighting conditions while looking at medium sized objects from a small distance, for example. On the one hand, we seem to have a fairly intuitive grasp of the concept: when I look out the window, I form my beliefs by eyesight, when I remember to bring the milk, I form my belief on the basis of memory. But these characterizations are somewhat open, we have a choice as to characterize them more specifically, for example as the method of forming beliefs by eyesight when it is day and lighting conditions are good, or more generally, such as the method of forming my belief on the basis of sensory impressions. As opponents of theories of justification referring to belief-forming methods have long been pointing out, there seems to be no objective basis of making a choice between these more general or specific characterizations. Without such a principled distinction, there seems to be no answer to the question what the process is that we use to form our belief. This problem is known as the generality problem (Conee and Feldman 1998; Goldman 1979). Besides the fact that this problem is not specific to safety-based approaches, and in fact plagues all major theories of epistemic justification, I will sidestep this problem by stipulation. For the argumentation theory literature contains many fairly specific descriptions of the various roles and steps in the process of argumentation (e.g. van Eemeren and Grootendorst 2016, Ch. 4), and we may simply select one of those processes and ask whether *that* process will help us eliminate luck.

to do so here. This specific approach has the virtues of specificity, recency and popularity, but the features of the approach that I will draw on in what follows, in particular about the roles of opponent and proponent, are generally accepted by other approaches, so our findings do not depend on the specifics of the pragma-dialectical theory.

As we saw above, in an unmixed dispute the role of proponent is to advance a thesis, and the role of opponent is to cast doubt on this thesis. Opponent may do so in different two ways; she may raise a motivated or an unmotivated challenge. A motivated challenge contains a reason for believing the standpoint is false, whereas an unmotivated challenge merely asks the proponent to provide additional support for the standpoint. Consider the following example:

P: It is raining outside

O: Why would that be true?

P: Because I see it raining through the window

O: But aren't experiences sometimes unreliable?

P: Well, lighting conditions are good and nothing out of the ordinary seems to be taking place...

Here, Proponent puts forward the thesis that it is raining outside, and Opponent tries to cast doubt on this thesis. First, by raising an unmotivated challenge, merely requesting reasons for the claim put forward. Second, by raising a motivated challenge that indicates a specific set of error-possibilities. I will now argue that in both cases, proponent's reasons function to eliminate error possibilities.

As I will understand reasons in this paper, they exclude ways in which the propositions for which they are reasons can be false. Thus, the claim that P sees it raining through the window is a reason for her claim that it is raining outside because the proposition that P sees it raining through the window is incompatible with certain scenarios where it is not raining outside—indeed, in this case, most normal scenarios; something strange would have to be going on for P to see it raining outside but it not actually raining. I take this to be a rather minimal requirement on reasons that expresses the fact that a reason for believing p is true is always at the same time a reason for believing p is not false. If a reason rules out

every scenario in which the target proposition is false, we say it provides deductive support for its conclusion, if it merely excludes some scenarios in which the target proposition can be false, the support provided is inductive.⁷ On this account, both deductive and inductive reasons exclude error-possibilities.⁸

What goes for the case of P and O goes for argumentation generally. If epistemic argumentation requires proponents to provide reasons for their claims, and if reasons are propositions whose truth rules out some possible ways in which the supported proposition could be false, then argumentative exchanges like the one above exclude error possibilities. Argumentation is thus a social belief-forming process that can eliminate error-possibilities.

Of course, arguments fail to provide support for their conclusions if their premises are false, and so the propositions advanced as reasons will only eliminate error-possibilities if they are *true*.⁹ My claim is not that all reasons eliminate error-possibilities, but only that they are able to do so. To support this claim, it is sufficient to show that true propositions can be reasons and that true propositions are incompatible with certain error possibilities.

If this much is admitted, then argumentation may contribute to the achievement of knowledge. As we have seen, in the course of argumentation specific error-possibilities may be eliminated. This contributes to the acquirement of knowledge because knowledge,

⁷ Note that this conception allows for very weak inductive reasons. Even if there are more p-worlds where *q* is false than there are p-worlds where *q* is true, *p* is still counted as being an inductive reason for *q* as long as *p* rules out some worlds in which *q* is false. However, this is compatible with the natural interpretation that in this case, while *p* provides some reason for believing *q* is true, *p* provides more reason for believing *q* is false. I thank an anonymous reviewer for pressing me on this point.

⁸ That still leaves the abductive case. This is a complicated case, since the nature and justification of abduction is contested (Douven 2017). While I do believe even abductive arguments exclude error-possibilities (after all, explaining *why* something is the case seems to involve explaining why it is not *not* the case) I leave the discussion of abductive arguments for another occasion, and restrict myself to deductive and inductive cases here.

⁹ This is an externalist conception of argument: an argument provides support only if it actually eliminates certain error-possibilities, not if we merely think it does. It follows that some arguments may be convincing even if they do not provide proper support.

as we saw above, requires that one's belief-forming process does not easily produce error. Because the process of argumentation as described in the pragma-dialectical theory leads to the exclusion of specific error-possibilities, proponent's belief in her claim will generally be safer after having gone through the process of argumentation than it was before. Again, the claim is not that all argumentation will necessarily have this result. But *when* argumentation excludes nearby error-possibilities, it contributes to the safety of arguer's beliefs, and so contributes to the acquirement of knowledge.

I want to stress again that my claim is not that argumentation always, or even predominantly leads to knowledge. The epistemic function may not be the primary function of argumentation, and additionally, the increase in safety of our beliefs may still be insufficient to amount to knowledge. But my aim in this section is to argue that argumentation can help us acquire knowledge, and for this an explanation of the way in which argumentation may increase the safety of our beliefs is sufficient. Argumentation is a tool that may legitimately be used for many purposes, my claim is merely that the epistemic purpose is among them.

It should be noted here as well that the pragma-dialectical framework that I have been using to argue for the claim that argumentation can produce knowledge is not itself an epistemic account of argumentation, in the sense that accounts of Goldman (2003) or Lumer (2005, 2020). According to such epistemic accounts, the primary function of argumentation is epistemic; usually the achievement of knowledge. As explained in Section 2, I adopt a pluralist account of (the aims of) argumentation. For my purposes, it is sufficient that argumentation *can* be used to achieve knowledge, and that is a claim that is compatible with both pragma-dialectical and purely epistemic accounts of argumentation. For even the pragma-dialecticians would allow that sometimes in resolving our differences of opinion, knowledge is generated.¹⁰

Seen in this light, the claim that argumentation can help us exclude error possibilities may seem fairly obvious. After all, what is argumentation other than a game of giving and asking for reasons, and what are reasons other than considerations that show a stand-

¹⁰ I thank an anonymous reviewer for pressing me on this.

point is true, i.e., *not false*? What is the value of specifically introducing the safety condition in this context? In general, I think it allows for a better, more detailed account of the relation between epistemic and dialectic norms. In particular, I will argue in the next section that our epistemic account can solve a well-known problem in argumentation theory.

5. The persistent interlocutor and epistemic perversity

In a well-known set of papers, Adam Leite and Michael Rescorla debate the question how to normatively assess the behaviour of a *Persistent Interlocutor* (PI). According to Leite, a persistent interlocutor is someone who, in argumentative contexts, does not cease to ask questions (2005, 397). In our terminology, a persistent interlocutor is an opponent who whenever proponent produces a reason in support of her claim, asks for further reasons supporting this claim. While asking for reasons is a legitimate move of opponent according to most theories of argumentation, we have the feeling that someone who *keeps* asking for reasons is doing something wrong, argumentatively speaking. The question is what this something is.

First, let us pin down opponent's unconditional right to ask for reasons, which Rescorla calls dialectical egalitarianism:

Dialectical Egalitarianism: all assertions require defence when faced with brute challenges (Rescorla 2009, p. 146).

Here a “brute” challenge is what we have been calling unmotivated challenges. According to Dialectical egalitarianism, proponent always has to respond to such challenges by providing reasons for her view.

The debate between Rescorla and Leite is primarily a debate about dialectical egalitarianism. Rescorla argues we should be egalitarians, while Leite argues we should instead be dialectical foundationalists, where dialectical foundationalism is simply the denial of egalitarianism; the claim that there are certain assertions that do *not* require defence when faced with brute challenges.

While the debate about the dialectical egalitarianism is a debate about the norms governing argumentation, it has potential epistemological consequences. One of the main reasons for foundationalism is that egalitarianism seems to lead to widespread scepticism. As the Pyrrhonians already saw, every assertion can be challenged by a persistent enough interlocutor. Since the Pyrrhonians also held that to be epistemically justified in holding a belief one needs to be able to successfully defend this belief in argumentative contexts, the impossibility of satisfying a persistent interlocutor means that the beliefs discussed in such contexts will not be justified. Since the persistent interlocutor's ability is content-independent, large swaths of our beliefs will turn out not to be justified. On a standard view where knowledge requires justification, widespread scepticism follows.

It is partly because of these dire epistemic consequences that foundationalists like Leite find room for propositions that can be asserted but cannot legitimately be challenged by the persistent interlocutor. Such a set can serve as the basis, or 'foundation' of all our other assertions, and we can defend all of our other beliefs successfully by drawing on them.

In other work I provide a more detailed analysis of the debate between Leite and Rescorla, and why I think that ultimately, both views are problematic (n.d). Here I am mainly interested in the exposition of my epistemic theory of argumentation, so I shall merely summarize the main points of critique. For Leite, I follow Rescorla in objecting that Leite's view unduly conflates epistemic and dialectical justification. Just because one is not able to provide an adequate defence of one's belief (dialectical justification), that does not mean it is not epistemically justified, as the many externalists about epistemic justification would agree. The main problem with Rescorla's view, however, is that it makes the behaviour of the persistent interlocutor epistemically irrelevant. On his view, the PI's behaviour does not violate any dialectical norm. She also, however, does not violate any epistemic norm. Rescorla is an epistemological foundationalist, so he thinks that our beliefs are ultimately justified on the basis of a set of foundational beliefs. On his view, the behaviour of the PI is epistemically irrelevant since nothing she says in conversation bears directly on this structure. In

terms of our example above: if proponent's belief that it is raining outside is properly based on her belief that she sees it raining outside, and this latter belief is properly foundational, then her belief is justified, no matter what opponent may say, and no matter whether proponent is then able to successfully defend her belief. As Leite puts the challenge to this view, however, we "want a way to say 'Shut up!' which isn't dogmatic but instead has the right sort of normative and justificatory force" (2005, p. 398). As I understand this comment, Leite means that we want to explain the specifically epistemic defects of the PI. Since for Rescorla the structure of dialectical and epistemic justification are disconnected, it is not immediately apparent how he could provide this explanation.

I believe our epistemic theory of argumentation does better in the sense that it evades both these criticisms. First, unlike Leite, the view does not maintain that epistemic justification requires dialectical justification. On the view explained above, in order to provide knowledge, the process of argumentation does not need to exclude *all* error-possibilities, only those that are metaphysically nearby. This means that the persistent interlocutor's request for reasons is only epistemically relevant if those reasons exclude nearby error. A request to eliminate far-off error-possibilities like the possibility that we are brains in vats or deceived by evil demons may be perfectly legitimate according to the norms of argumentation, but since knowledge does not require these error-possibilities to be eliminated, such argumentation will be epistemically *irrelevant*. Insofar as we use argumentation to acquire knowledge, that is, insofar as we are concerned with *epistemic* argumentation, responding to the persistent interlocutor is not necessary and we may acquire knowledge even without an answer to all brute challenges. So, on our view, successful epistemic justification does not require successful defence against the persistent interlocutor. Scepticism is averted even without reliance on dialectically basic propositions.

On the other hand, the present view quite naturally explains how we may criticize the persistent interlocutor on an epistemic basis. While the persistent interlocutor does not violate any dialectical or epistemic norms, she does frustrate the epistemic goal of argumentation. We sometimes do use argumentation to acquire

knowledge as I have explained above, and if the persistent interlocutor in these contexts requests reasons that would eliminate far-off error possibilities, she asks the participants to the debate to spend cognitive effort on an activity that is epistemically irrelevant. We cannot, but also need not answer such challenges in order to acquire knowledge.¹¹

This criticism of the persistent interlocutor is partly pragmatic in character: I have argued that her behaviour is problematic because it frustrates some of the aims with which we may engage in argumentation. It is however, also partly epistemic because the aim that is frustrated is the aim of acquiring knowledge. I believe this view provides a better explanation of the epistemic dimension of the persistent interlocutor than Rescorla's view; better because unlike Rescorla, the present view is able to explain how the persistent interlocutor may hinder our quest for knowledge even while admitting that dialectical and epistemic justification are distinct.

All of this is not to deny that argumentation may serve other purposes than that of producing knowledge. For such other purposes, raising far-off error-possibilities in argumentation may be beneficial. But from the epistemic point of view, there is something amiss with discussing such issues; this discussion takes up valuable cognitive resources and contributes nothing to the acquirement of knowledge. Accordingly, we shall call argumentation that involves the discussion of far-off error-possibilities *epistemically perverse*, in contrast to epistemically legitimate argumentation that involves nearby error-possibilities.

¹¹ An anonymous reviewer has objected that my criteria rules out all debates on fundamental issues, since questioning our fundamental premises tends to involve far-off error possibilities rather than nearby ones. In response, I would like to emphasize that my criterion is an externalist one. The world is a certain way, and given that way, certain errors are easily possible, and others are not. If our most central beliefs about the world are false, then on my account it is perfectly legitimate to discuss them since they are false in very near worlds. But if they are not, then discussing them is epistemically perverse. So, it would be fine on my account for medieval scientists to question the assumption that the earth was flat, but it is (presumably) epistemically perverse for current flat-earthers to question that it is round. Thanks to the anonymous reviewer for pressing me on this point.

6. Three objections

Let us take stock. So far, I have argued for a safety-based account of epistemic argumentation. This account of argumentation explains how argumentation may serve to produce knowledge. In the previous section, I explained how the account can provide a new and better perspective of a notorious character in argumentation theory. In this section I will complete the defence of my account by considering some potential objections.

I first discuss two general objections to epistemic theories of argument, drawing on the excellent exposition of Scott Aikin (2008).¹² The first objection runs as follows:

because epistemological theories are highly controversial within epistemology, it is best to get as much as one can for argumentation theory independently of the controversial theories. Having one's entire case for an epistemic theory of argument hang on a highly controversial premise in, say, meta-epistemology is recipe for a theoretical disaster. (Aikin 2008, p. 140)

This objection may be run against our theory: in my defence of a modal account of epistemic argumentation, I have relied on the principle of safety. If this principle is highly contested, my theory will inherit this contestability.

Luckily for us, the safety condition is not highly contested. While historically, some alternative modal proposals have been made, such as Robert Nozick's sensitivity condition (1981), it seems clear by now that the safety principle offers a superior formulation of the modal profile of knowledge (Sosa 1999). One of the main reasons for this is that virtually everyone agrees that Gettier cases fail to constitute knowledge due to the presence of certain kinds of knowledge-undermining luck, and that the safety principle

¹² Aikin has his own solutions to these problems. They resemble the present approach in that Aikin recognizes a potential pluralism of the norms governing arguments depending on the ends with which the arguments are used. My approach thus falls in the same category of views as that of Aikin. It differs from Aikin's approach, however, in that I focus on the externalist epistemic norm of safety, while Aikin remains uncommitted as to which epistemic norm applies to argumentation, but seems in general more sympathetic towards internalist norms (see my response to the second objection below).

is specifically designed to exclude this kind of luck (Pritchard 2005). So, while there are disagreements over the precise formulation of the safety principle, and over the question whether knowledge requires *more* than just safe belief, the fact that knowledge requires some form of safe belief is relatively uncontroversial.

While in this paper I have opted for a specific formulation of the safety condition to be able to show how argumentation may increase the safety of our beliefs, I believe what I have said can easily accommodate slightly different formulations of the safety principle. The main point I have made in this paper is that focusing on the modal requirements for knowledge supports the epistemic theory of argumentation because it explains how argumentation may produce knowledge, and sheds new light on the legitimacy of the persistent interlocutor's behaviour.

Furthermore, it bears emphasizing that my general view of argumentation is pluralistic: we can argue with different aims, and relative to those different aims different norms will apply. In this paper, I talk only of epistemic argumentation: argumentation with the aim of acquiring knowledge. That we use epistemic principles to shed light on this kind of argumentation seems necessary: what will be good norms for acquiring knowledge through argumentation will very much depend on what knowledge is. Of course, the general analysis of knowledge is controversial. But I hope to have shown that we can get quite a lot of theoretical clarification out of focussing on less controversial features of knowledge like the safety condition. That we may not have a complete definition of knowledge means that we will not be able to specify the norms for epistemic argumentation exhaustively.¹³ But it does not mean that we cannot specify any norm for epistemic argumentation.

The second objection that I want to discuss also starts from the observation that my story depends on a particular epistemic norm. Contrary to the foregoing critique, however, it takes issue with the *externalist* nature of my safety condition specifically. Again, we can refer to Aikin's specification of the objection:

¹³ Although see de Grefte, J. (2021) for an attempt at a definition.

The practicability objection runs that epistemic theories are either too thin to answer the needs of offering advice or offer useless advice, because epistemic theories risk ignoring the attitudes of those who need correction. (Aikin 2008, p. 134)

The complaint here is that using epistemic theories as criteria for good argument is a bad idea because such criteria will not be able to provide advice that arguers can follow, since the epistemic theories are often thin (so that it is hard to draw practical consequences from them) and moreover, in order to have any consequences in the first place, these principles need to be recognized by the participants as such, which may not often be the case, given the narrow class of people with epistemic knowledge.

My first response to this objection is that safety accounts of knowledge are not too thin to draw practical consequences from. Aikin's example of a problematically thin notion is that it would only be able to provide advice like "construct arguments that provide good epistemic reasons" (Aikin 2008, p. 132). Without any specification of what a good epistemic reasons are, this advice is fairly useless. However, it is exactly the notion of good epistemic reasons that is given some clear content by a safety-based approach. Good epistemic reasons are reasons that exclude nearby error-possibilities in the sense described above. Since we can determine which reasons exclude which error-possibilities, this is advice that is practically followable; we aim for reasons that exclude as much error as possible, until the error-possibilities excluded are clearly far-off, and therefore irrelevant to our epistemic pursuit.

Of course, we may in principle be wrong about what the actual world is like, and we may disagree about which worlds are similar enough to the actual one to be relevant for knowledge. But at least we know what we are looking for. We know at least roughly what it would take to rule out nearby error. That the boundary is vague between nearby and not-nearby error means that it may not always be possible to determine whether a particular challenge raised by PI is problematic. But that does not prevent us from criticizing the PI if she is raising clearly far-off error possibilities. Even if we cannot always be certain that we currently satisfy the criterion, the criterion is not devoid of content and can thus be applied in argumentative contexts.

My second response is twofold. First, even if arguers are not explicitly aware of many epistemic principles, they may still be implicitly aware of them. I do not take the present account to be revisionary, in the sense that I believe that most arguers most of the time already try to exclude easily possible error. As I have argued above, reasons can exclude error-possibilities, whether people providing these reasons are aware of this feature or not. As long as the reasons they provide exclude nearby error, these practices provide knowledge, whether the people involved in argumentation are aware of these norms or not. Secondly, I want to stress that I am concerned in this paper with *epistemic argumentation*, that is argumentation directed at the acquirement of knowledge particularly, not argumentation in general. And follows from the aim of this kind of argumentation that the norms on knowledge apply to this kind of argumentation as well, whether participants are aware of these norms or not. Consider the fact that knowledge requires truth. Participants may not be aware of the fact that the beliefs produced by their argumentation are true or not, but argumentation that produces false belief is defective *qua* epistemic argumentation. Such argumentation does not fulfil its goal, and it can thus be criticized *as* epistemic argumentation. Similarly, argumentation that does not produce safe belief can be criticized as epistemic argumentation, whether the participants to the discussion have any means of determining the safety of their arguments or not.¹⁴

The final objection I want to discuss is specific to approaches utilizing safety conditions. Some may object to our condition on

¹⁴ Another way to make the same point is in terms of the internalist vs externalist distinction in epistemology. Especially if we include truth as a criterion on knowledge, it is fairly uncontroversial that knowledge is an externalist notion in the sense that whether we know is something that does not supervene on our mental states or that we can establish by reflection alone. If knowledge is an externalist notion, and epistemic argumentation aims at knowledge, then we will not be able to determine by reflection alone whether our argumentation produces knowledge. Still, this is crucial for whether our argumentation meets its aims, and so for whether our argumentation meets the norms of epistemic argumentation. It follows that the norms of epistemic argumentation are at least partly externalist norms.

epistemic argumentation that the notion of safety is not precise enough to serve in a fruitful clarification. While we may have an intuitive grasp of notions like world similarity and belief-forming methods, these notions are notoriously hard to make precise.

As we have been arguing so far, a belief is safe if and only if it is produced by a belief-forming method that would not easily produce false belief. This means our analysis of (epistemic) safety crucially rests on the notions of belief-forming method on the one hand and the notion of easy possibility (spelled out in terms of worlds similar to the actual world) on the other. Our current objector argues that these notions are too vague to serve their purpose. In response, I will look at both notions in turn and argue that while they are vague, this does not detract from their theoretical usefulness.

First, consider the notion of belief-forming method. The objection to using the notion in an epistemic analysis is not necessarily that we cannot define belief-forming methods in a precise way, for we may say that belief-forming methods are processes that take inputs of type A and output beliefs of type B, and fill in A with visual percepts and B with beliefs about the external world, and then the resulting analysis is relatively precise. The charge of vagueness here rather concerns that it is indeterminate what prespecification we should choose on any occasion. Consider this example: I look at a tree and form the belief that there is a tree there. We may fill in for A the following:

- 1) Visual percepts
- 2) Visual percepts with tree like appearances
- 3) Visual percepts with tree like appearances in good lighting conditions
- 4) Visual percepts with tree like appearances in good lighting conditions on April 19, 2021.

These are various correct descriptions of the inputs of my belief-forming method. We similarly have choice on how to fill in B in our schema; we may put in ‘beliefs,’ ‘beliefs about the external world,’ or ‘beliefs about trees,’ for example. Crucially, the safety of the method may differ under these various descriptions: it is plausible that forming any beliefs on the basis of any visual percepts may more easily go wrong than the formation of tree beliefs

on the basis of visual percepts with tree-like appearances under good lighting conditions. In our possible worlds framework, this can be explained in the following way. When we check for safety; we have to keep the method fixed and check whether there are nearby worlds where that method is used but it produces a false belief. If we define the method very broadly, then more nearby worlds may be relevant and so the probability for false beliefs in nearby worlds increases if we define our methods broadly.

Of course, the objection considered here is a version of the well-known ‘generality problem’ for reliabilism. This makes my defence somewhat easier, because it is now not exclusively on my shoulders to find a solution to this problem. In fact, as several authors have argued recently, the generality problem is a problem for a wide range of distinct epistemologies (e.g. Bishop 2010; Comesaña 2006). While one can take this as so much the worse for those epistemologies, the fact that many theories from different epistemological orientations all run into some form of the problem means that it may be inescapable. I suspect it is, since it seems clear that whether we know depends on the way we have formed our beliefs. The same belief that will never amount to knowledge when we arrive at it through tea-leaf reading may nevertheless constitute knowledge if formed in a different way, say by consulting an encyclopaedia. Reflection on cases like this suggests that the difference between knowledge and merely true belief is constituted at least in part by the method of belief-formation. If this is true, then no matter what exact analysis is given, any plausible epistemological theory will need to refer, in some way, to belief-forming methods, processes or ways in which we form our beliefs. The generality problem is an important epistemological problem to be solved, but not one that counts against any particular epistemological theory.

While I do not aim to solve the problem here, I think a potential dissolution to the problem may actually be suggested by looking at various non-epistemic forms of safety. First note that for these other kinds of safety, a similar problem arises: crossing the street may be relatively safe in general, but not when we look primarily at our cell-phone, listening to loud music. How we describe the way we do things matters for their degree of safety. While the

generality problem is thus a potential problem even outside epistemology, the fact that we have no principled way of describing our methods of action has not, and should not have prevented us from improving the safety of those methods. Looking left-right-left when crossing the street increases our safety, whether we have a principled way of describing the method or not. The hope implicit in a safety-based epistemology is that once it is clear that we aim to increase the safety of our beliefs, it will be able to similarly identify pieces of conduct in epistemology to do so. Like the example of looking left-right-left, the prescribed behaviour may be fairly particular, and only increase the safety in specific situations (in left-driving countries it will not produce the same degree of safety). But we may also hope to uncover more general methods that help us increase the safety of our beliefs in a wide range of circumstances. As I argued above, argumentation may be just such a method.

A second source of vagueness may come from the notion of world similarity. As indicated above, our notion of safety depends on how similar the nearest possible world is where the target event occurs. While the notion of similarity is itself a notion that has resisted philosophical analysis, the notion of complete world similarity may seem even more difficult to make any sort of precise sense of.

Yet, the notion has been fruitfully applied in philosophy of language in the analysis of counterfactuals, among other things (Lewis 1973; Stalnaker 1968). And also, our psychological capacity for counterfactual thinking, so central to many of our reasoning and planning abilities (Williamson 2005), for example, seems to depend on *some* way of ranking possibilities in terms of how easily they could have occurred.¹⁵

We thus seem to have an intuitive, coarse grained method of answering questions like which worlds are more easily possible

¹⁵ This is not exactly the same capacity as the capacity for probabilistic reasoning. Consider playing Russian roulette with one bullet in a thousand chambers: probabilistically, in this situation one does not easily lose. There seems to be a clearly different sense of easy possibility, however, in which in this situation it *is* easily possible that one shoots oneself.

than others. The coarse and intuitive nature of this method means that there will be borderline cases, where it is not clear whether a world where our belief is false is similar enough to count as easily possible. But it is an open question whether this is problematic in the context in which the notion of safety is used in epistemology. If it is not clear whether our belief-forming method may have easily produced a false belief, it may be equally unclear whether the resulting belief constitutes knowledge. That there is no exact degree of safety required for knowledge seems a fact to be explained rather than an objection to an epistemic theory. Forming beliefs about medium-sized objects in utter darkness is too unsafe a method to preclude knowledge—that much seems clear—but looking at the same objects in broad daylight seems safe enough. That there are going to be cases in between where it is just not clear whether we acquire knowledge about these objects just by looking seems fairly obvious. One upshot of the present paper is that the same goes for the method of argumentation: sometimes the support provided will clearly be strong enough for producing knowledge, sometimes it will clearly not be strong enough, and sometimes, it will be unclear whether an argument is strong enough to provide knowledge.

The fact that it is sometimes vague whether arguments provide knowledge does not detract from the usefulness of the epistemic theory of argument. Insofar as epistemic argumentation is concerned, arguments are legitimate only if they are safe enough to produce knowledge. If it is unclear whether this condition is satisfied, it will be sufficient to provide further support until it *is* clear that the resulting beliefs, if true, would constitute knowledge. Of course, we cannot exclude all error-possibilities, but the whole point of a safety-based analysis of knowledge is that this is not required: we need to exclude only those error-possibilities that are easily possible, and this is usually within our limits.

7. Conclusions

In this paper, I first outlined an inclusive theory of argumentation. I then focussed on epistemic argumentation specifically, and proposed a safety-based norm for this kind of argumentation. The

approach is more modest than extant epistemic theories of argumentation by restricting the normative force of the epistemic norms on argumentation to epistemic argumentation specifically: if we do not aim to argue for knowledge, then we are not bound by the epistemic norms, but once we try to achieve knowledge through argument, the norms take effect. I have spelled out in some detail what I take to be the most important epistemic norm on argumentation, the safety condition. Since current epistemic approaches to argumentation are usually spelled out in reliabilist or evidentialist terms, a safety-based approach is novel. I have explained the best way to cash out the notion of safety, and shown how the notion may be fruitfully applied to shed light on the epistemic and argument-theoretic validity of the behaviour of the notorious character of the persistent interlocutor. Finally, I have responded to what I take to be the most obvious objections to the present theory. While a full defence of the safety-based approach goes well beyond this paper, I hope to have shown that it is a serious contender in the ring of epistemic theories of argument, and provides a fruitful addition to argumentation theory generally.

References

- Aikin, Scott F. 2008. Three objections to the epistemic theory of argument rebutted. *Argumentation and Advocacy*. 44(Winter): 130-42.
- Bishop, Michael A. 2010. Why the generality problem is everybody's problem. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 151(2): 28-98.
<<http://www.jstor.org/stable/40926824>>
- Comesaña, Juan. 2006. A well-founded solution to the generality problem. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 129(1): 27-47.
- Conee, Earl and Richard Feldman. 1998. The generality problem for reliabilism. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 89(1): 1-29.
- Douven, Igor. 2017. Abduction. *Stanford Encyclopedia of Philosophy*, 1-22. Metaphysics Research Lab, Stanford University.
- Dutilh Novaes, Catarina. 2013. A dialogical account of deductive reasoning as a case study for how culture shapes cognition. *Journal of Cognition and Culture*. 13(5): 459-82.

- Dutilh Novaes, Catarina. 2021. Argument and argumentation. In *Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. Metaphysics Research Lab, Stanford University. Fall 2021 ed.
- Eemeren, Frans H. van and Rob Grootendorst. 2003. *A systematic theory of argumentation: The pragma-dialectical approach*. Cambridge: Cambridge University Press.
- Eemeren, Frans H. van, and Rob Grootendorst. 2016. *Argumentation, communication, and fallacies: A pragma-dialectical perspective*. Routledge.
- Eemeren, Frans H. van and Rob Grootendorst. 2004. *A systematic theory of argumentation: The pragma-dialectical approach*. Cambridge University Press.
- Goldman, Alvin I. 1979. What is justified belief? In *Justification and knowledge*, ed. G. S. Pappas. Dordrecht: D. Reidel Publishing Company.
- Goldman, Alvin I. 2003. An epistemological approach to argumentation. *Informal Logic*. 23(1): 51-63.
- de Grefte, J. 2020. Towards a hybrid account of luck. *Pacific Philosophical Quarterly* 101(2), 240-255.
- de Grefte, J. n.d. The persistent interlocutor.
- de Grefte, J. 2021. Knowledge as justified true belief. *Erkenntnis*, 1–19.
- Leite, Adam. 2005. A localist solution to the regress of epistemic justification. *Australasian Journal of Philosophy*. 83(3): 395-421.
- Lewis, David K. 1973. *Counterfactuals*. Cambridge, MA: Harvard University Press.
- Lumer, Christoph. 2005. The epistemological theory of argument--how and why? *Informal Logic*. 25(3): 213-243.
- Lumer, Christoph. 2005. Introduction: The epistemological approach to argumentation--A map. *Informal Logic*. 25(3):189-212.
- Lumer, Christoph. 2020. Knowledge and the epistemic function of argumentation. In *Reason to Dissent. Proceedings of the 3rd European Conference on Argumentation*, eds. Catarina Dutilh Novaes, Henrike Jansen, Jan Albert van Laar and Bart Verheij, 219-224. College Publications.
- Nozick, R. 1981. *Philosophical explanations*. Cambridge, Massachusetts: Harvard University Press.
- O'Keefe, Daniel J. 1992. Two concepts of argument. In *Readings in argumentation*, eds. William L Benoit, Dale Hamble and Pamela J Benoit. Foris Publications.
- Perelman, Chaïm and Lucie Olbrechts-Tyteca. 1969. *The new rhetoric: A treatise on argumentation*. Notre Dame, University of Notre Dame Press.

- Pritchard, Duncan. 2005. *Epistemic luck*. New York, NY: Oxford University Press.
- Pritchard, Duncan. 2012. Anti-luck virtue epistemology. *The Journal of Philosophy* 109(3): 247-79.
- Rescorla, Michael. 2009. Epistemic and dialectical regress. *Australasian Journal of Philosophy* 87(1): 43-60.
- Sosa, Ernest. 1999. How to defeat opposition to Moore. *Nous* 33(s13):141-53.
- Stalnaker, Robert C. 1968. A theory of conditionals. In *Studies in logical theory*, vol. 2, ed. Nicholas Rescher. Oxford: Blackwell.
- Williamson, Timothy. 2000. *Knowledge and its limits*. Oxford, NY: Oxford University Press.
- Williamson, Timothy. 2005. Armchair philosophy, metaphysical modality and counterfactual thinking. *Proceedings from the Aristotelian Society* 105(1): 1-23.
- Williamson, Timothy. 2009. Probability and danger. *The Amherst Lecture in Philosophy* 4: 1-35.
<<http://www.amherstlecture.org/williamson2009/>>