

Understanding echo chambers

André Jonassen

Masteroppgave i filosofi

Veiledet av professor Olav Gjelsvik

Institutt for filosofi, ide- og kunsthistorie og klassiske språk
Høsten 2022



Forord

Jeg ønsker å takke min veileder Olav Gjelsvik for all tålmodighet og hjelp med oppgaven. Han har vært helt avgjørende for å finne oppgavens retning. Videre vil jeg takke mine foreldre for god støtte i innspurten og uvurderlig hjelp med gjennomlesning av og diskusjoner om oppgaven. Jeg vil også takke min samboer Anna som, til tross for ha feilaktig blitt redusert til en høyreekstrem og epistemisk uinteressert person i et av denne oppgavens eksempler, har vært veldig snill, grei og tålmodig med undertegnede under hele skriveprosessen. Til slutt vil jeg også takke hunden Chomsky som kanskje ikke kommer med like mye lur innsikt som sin akademiske navnebror, men som likevel har vært en flott samtalepartner under denne oppgavens innspurtsfase.

Abstract

Some recent definitions of echo chambers hold that echo chambers are “social epistemic structure[s] in which other relevant voices have been actively discredited” (Thi Nguyen, 2020, 142). This definition does a good job at describing *what* echo chambers are, but does not go equally far in describing *why* echo chambers are, that is, which type of behavior enables them to exist. Further, without a stronger focus on why echo chambers engage in exclusion and discreditation of relevant voices, this definition also covers social epistemic structures which are in no way epistemically harmful. In other words, with this definition, we can call epistemic communities that do not necessarily obstruct their chances of gaining knowledge through their exclusion and discreditation echo chambers. I argue that a definition of echo chambers should be focused on epistemic structures that are epistemically harmful, that damages the echo chambers members’ prospects of gaining knowledge. Therefore, Thi Nguyen’s definition can be revised in order to make it focused. My revised term suggests that we can connect the concept of epistemic vices in order to explain why members of echo chambers actively discredits and excludes relevant voices. Epistemic vices are “[...] personal intellectual failings that have a negative impact on our intellectual conduct” (Cassam, 2019, 2), some examples being prejudice, gullibility and closed-mindedness. I suggest that we define echo chambers as social epistemic structures in which relevant voices are actively discredited *due* to one or more such epistemic vices. With this definition, one limits the terms scope to those social epistemic structures which are justified in epistemically vicious behavior, that is, those structures whose very foundation are grounded in behavior that gets in the way for gaining knowledge. Limiting the scope to only include such structures has to primary effects. Firstly, it provides us with a framework to understand why non-vicious structures in fact are not echo chambers, and by effect makes our inquiries regarding echo chambers focused on those that are epistemically vicious. Secondly, it enables us to focus our attention to the behavior that enables echo chambers to exist, a focus I believe to be vital in any attempt to find effective measures to combat echo chambers.

Table of contents

- Forord 2**
- Abstract 3**
- Introduction 5**
 - Theme and background..... 5*
 - Topic and terms..... 6*
 - Question 7*
 - Motivation..... 7*
 - Thesis statement 8*
 - Methodological approach 8*
 - Limitations..... 10*
 - State of the art 10*
- Thi Nguyen’s Echo Chambers and Epistemic Bubbles 11**
 - Exiting an epistemic bubble..... 14*
 - Exiting an echo chamber 16*
- Epistemic vices 22**
 - Stealthy vices..... 27*
 - Obstacles to vice detection..... 28*
 - Overcoming the obstacles to vice detection..... 29*
 - Prospects of self-improvement..... 32*
- Epistemic vices and echo chambers 34**
 - Epistemically vicious echo chambers..... 34*
 - Understanding Thi Nguyen’s proposed exit strategy through the vice-lens..... 36*
- Are all echo chambers malignant? 41**
 - Academics as an echo chamber 41*
 - Academics as seen through the vice-lens 45*
- Consequences of the existence of non-malignant echo chambers..... 49**
 - A revised definition of echo chambers 50*
- Conclusion..... 54**
 - Further research 56*
- Literature 57**

Introduction

Theme and background

Many of us get our information from just a limited number of sources. We might simply not surround ourselves with enough people, or have a fondness for just one single news network, or have a social media algorithm that focuses our attention just to a few voices. Moreover, and maybe more importantly, some of us seem to increasingly expose themselves to sources that all are in some kind of agreement about some fundamental set of facts. Being exposed to just a few sources and those sources being in agreement on some facts is not automatically a bad thing, but it can become quite challenging if those accepted facts are far from accepted by those exposed to another set of sources.

One of the potential challenges that can be caused by people limiting themselves to considering sources that supporting a specific angle or a specific set of beliefs that do not conform with alternative sources is that their beliefs become more extreme and less grounded in reality. This is because the information they receive is not balanced or objective, and is instead tailored to support their existing beliefs. As a result, their views may become more extreme and less nuanced, and they may be less open to considering other perspectives. This process is often what referred to as an echo chamber effect.

Furthermore, the echo chamber effect can lead to a lack of understanding and empathy for people with different beliefs. When people only engage with sources that align with their own beliefs, they are not exposed to opposing viewpoints, which can lead to a lack of understanding and empathy for those who hold different views. This can ultimately fuel conflict and polarization, as people become more entrenched in their beliefs and less able to see things from other people's perspectives.

Additionally, the echo chamber effect can limit people's access to potentially important information. When people only engage with sources that support their existing beliefs, they may miss out on information that could be relevant or valuable to them. This can lead to a lack of knowledge and understanding on a wide range of topics, and can make it difficult for people to make well-informed decisions.

Overall, the echo chamber effect can have a number of negative consequences, including making people's beliefs more extreme and less grounded in reality, decreasing understanding and empathy for people with different beliefs, and limiting access to important information. It is therefore important for people to be aware of the echo chamber effect and to make an effort to seek out diverse sources of information, in order to avoid becoming trapped in an echo chamber.

Topic and terms

The main topic of this paper is already been mentioned several times, namely echo chambers. Echo chambers are often understood as communities with a limited intake of information, and where the lack of external information amplifies the internal voices. This means that the people within the echo chamber are only exposed to a narrow range of viewpoints and perspectives, which can lead to the reinforcement of their existing beliefs and a lack of critical thinking.

Echo chambers may also sometimes used in the context of social media and social media algorithms. In this context, echo chambers refer to the way that social media algorithms can filter and curate the information that users see, based on their previous activity and interactions. This can potentially create a structure in which you are only exposed to beliefs that conforms with your pre-existing beliefs, limiting your exposure countering arguments and views.

In this thesis we will focus on the definition of echo chambers presented in N. Thi Nguyen's paper *Echo chambers and epistemic bubbles* (2020) as our starting point. He defines echo chambers as “[...] a social epistemic structure in which other relevant voices have been actively discredited [...]” (Thi Nguyen, 2020, 142)

The subtopic of this paper is *epistemic vices*. As presented in *Vices of the mind* (2019), epistemic vices are “[...] personal intellectual failings that have a negative impact on our intellectual conduct” (Cassam, 2019, 2). Examples of epistemic vices are overconfidence, prejudice, and closed-mindedness, since they all are failings of the mind that get in the way of intellectual conduct such as gaining knowledge.

Another central topic of this paper is the concept of epistemic vices. As defined by Quassim Cassam in *Vices of the Mind* (2019), epistemic vices are "personal intellectual failings that have a negative impact on our intellectual conduct" (Cassam, 2019, 2). In other words, epistemic vices are habits or dispositions that can negatively affect our ability to acquire, evaluate, and use knowledge. Examples of epistemic vices include overconfidence, prejudice, and closed-mindedness, as these tendencies can interfere with our intellectual conduct and prevent us from gaining new knowledge.

We will go in much more depth about both echo chambers and epistemic vices throughout the thesis.

Question

In this thesis, I wish to explore the concept of echo chambers. My aim is to consider definitions of echo chambers, their precondition and which epistemic mechanisms that are in motion within the echo chamber members. The research question will thus be the following: what are echo chambers and what enables them to exist?

Firstly, I wish to distinguish echo chambers from similar epistemic structures such as epistemic bubbles and get a more to the point conception of the concept. Secondly, I wish to explain the epistemic mechanisms which take place within members of echo chambers in terms of epistemic vices. While I believe Thi Nguyen's paper on echo chambers is a good starting point for understanding *what* echo chambers are, I argue that considering echo chambers through the lens of epistemic vices can help us understand *why* they are. Thirdly, I wish to make the case that echo chambers understood as Thi Nguyen does are not necessarily malignant.

Motivation

The motivation for these topics lie in an interest in the broader phenomena of post-truth and fake news. These phenomena seem to have become increasingly prominent in recent years, as more and more people have become skeptical of the information they receive from traditional sources, such as the news media, government institutions, and experts. This has led to the rise of alternative sources of information, such as social media and closed online forums, which

often spread misinformation, conspiracy theories, and other types of false or misleading information.

So, the project first has a theoretical motivation in understanding some of these contemporary phenomena, in seeking answers to questions like how echo chambers come into being, how they function, and which epistemic processes they rely on. This theoretical motivation leads to the practical motivation of resolve: how does one break out of an echo chamber, or how can one minimize the risk of echo chambers coming to being in the first place?

Answering some of these questions can, hopefully, serve as a step towards a more unified public debate, where one at the very least can base one's argumentation on the same set of factual premises regardless of which side of the aisle you find yourself in.

Thesis statement

It will be fruitful to narrow a definition of echo chambers to mean not only a social epistemic structure in which other relevant voices have been actively discredited, as Thi Nguyen holds, but additionally that this discreditation is justified by one or more epistemic vices.

Understanding echo chambers in terms of epistemic vices helps us better understand why epistemic agents partake in them, because one can then shift one's focus towards identifying the underlying epistemic vices causing the membership to begin with. Echo chamber members can be guilty of several epistemic vices and understanding which vices echo chamber members are guilty of makes it easier (although not at all easy) to more precisely identify the "original sin". A vice-view on echo chamber also excludes social epistemic structures that are *not* epistemically vicious from being echo chambers.

Methodological approach

In the following, I will briefly outline the approach taken in this thesis.

Firstly, we will examine Thi Nguyen's paper on echo chambers and epistemic bubbles in order to get a clear terminological and conceptual point of departure. After our inquiry of Thi Nguyen's theory, we will have a temporarily clear conception of what echo chambers are, as

well as what they are not. Further, we will examine Thi Nguyen's proposed exit-strategy, so to have an indication of how an echo chamber member can break out. This exit-strategy will be re-examined later in the thesis.

Secondly, we will introduce Cassam's theory on epistemic vices. We will first get a general conception of epistemic vices, before we engage in a more detailed inquiry of what Cassam calls "stealthy vices". This inquiry includes examining reasons for why detecting one's own stealthy vices is much more difficult than detecting one's non-stealthy vices, as well as seeing how Cassam proposes one can overcome these difficulties. Furthermore, we will see how an epistemic agent can go from detecting a vice to actually doing something about it when we consider Cassam's discussion on the possibility of self-improvement.

Thirdly, we will return to considering echo chambers following Thi Nguyen's definition, but this time the consideration will revolve around epistemic virtues. We will consider how echo chambers function and what enables them to exist by utilizing Cassam's theory. A large part of this re-examination includes attempting to explain Thi Nguyen's proposed exit strategy by utilizing the concept of epistemic vices. We will here re-examine an example used by Thi Nguyen and see how a single echo chamber member's behavior can be explained through epistemic vices, covering a process going from echo chamber membership to leaving that echo chamber.

Resulting from this vice-view of echo chambers, we will see that although many echo chambers are epistemically vicious, echo chambers are not *necessarily* epistemically vicious. This leads us to a discussion on how a non-vicious echo chamber may look like. We will go in depth explaining how a possible non-vicious echo chamber could function, and we will see exactly why it is distinct from vicious echo chambers.

We will then, concludingly, draw motivation from the distinction between vicious and non-vicious echo chambers to create a revision of Thi Nguyen's definition of echo chambers. The newly revised term will be presented and explained, before we again revisit some of the examples used in this thesis in order to understand how the revised term changes how we make sense of different social epistemic structures.

Limitations

It is important to recognize that this thesis has certain limitations, and that there are certain aspects of echo chambers that it does not attempt to address. One such limitation is the focus on the individual epistemic behavior of members of an echo chamber, rather than on the collective dynamics of the echo chamber itself. While this approach allows us to gain a better understanding of the ways in which individual agents engage with echo chambers, it does not fully capture the social epistemic structures and processes that underlie echo chambers. This is an important area for future research, as the collective dynamics of echo chambers can have implications of great importance regarding how echo chambers are formed and maintained, as well as how echo chambers affect their members. Examining how echo chambers by examining it from a social epistemic viewpoint can grant us a more complete understanding of how echo chambers operate and can potentially help us develop effective strategies for mitigating their negative effects.

State of the art

In contemporary epistemologically, there seems to be a rise in authors discussing social epistemology generally and echo chambers specifically. In the next, I will briefly introduce both the texts central in this paper as well as some that go beyond the paper's scope while remaining important in the context of echo chambers.

The two works that this thesis will build upon have both already been mentioned, namely C. Thi Nguyen's *Echo chambers and epistemic bubbles* (2020) and Quassim Cassam's *Vices of the Mind* (2019).

In his paper, C. Thi Nguyen distinguishes between epistemic bubbles and echo chambers, the difference being that in the former one lacks exposure to relevant information and arguments, in the latter one systematically distrusts outside sources (Thi Nguyen, 141, 2020). So, per Nguyen, an epistemic bubble can burst in a meeting with new arguments, while an echo chamber might only be reinforced by counterarguments.

In *Vices of the Mind*, Quassim Cassam develops an incredibly thorough account of epistemic vices. Building his theory with basis in quite recent events such as the American presidential election of 2016 and the vote on Brexit, he discusses what epistemic vices are, the damage

they can do, why they are so incredibly difficult to get rid of, and, importantly, how there actually is something we can do about them despite all the difficulties one will meet in attempting to do so.

There is, of course, much relevant material that does not get covered in this thesis. I would like to mention two of these. What they both have in common is their discussions of epistemic structures as groups, and I think both works can contribute in moving beyond the scope of my thesis, shifting the focus from the individual to the collective.

Kenneth Boyd builds on Thi Nguyen in his paper *Epistemically Pernicious Groups and the Groupstrapping Problem* (2019). According to Boyd, epistemic bubbles are characterized by receiving information from a limited number of sources, those sources often supporting your already held beliefs. Echo chambers function in the same way, but additionally actively reinforce certain beliefs (Boyd, 61, 2019). Further, Boyd moves toward an analysis of groups rather than individuals and introduces what he calls ‘the groupstrapping problem’ as one explanation for why epistemic bubbles and echo chambers are hard to break out of.

John Greco provides a fundament for theories on echo chambers and epistemic bubbles in his book *The transmission of knowledge* (2020). In it, Greco creates a thorough account of testimonial knowledge, claiming that testimony is a social act governed by social settings and norms. Understanding the

Thi Nguyen’s Echo Chambers and Epistemic Bubbles

In this section I will present C. Thi Nguyen's article "Echo chambers and epistemic bubbles" to explain the notions of "epistemic bubbles" and "echo chambers", and then go on to discuss what is epistemically problematic about each of them using concrete examples. I will also discuss his account of how it is possible to escape from an echo chamber, briefly start discussing a possible Cartesian way of escaping them.

Thi Nguyen defines epistemic bubbles as the following: (Thi Nguyen, 2020, 142)

[A]n epistemic bubble is a social epistemic structure in which some relevant voices have been excluded through omission. Epistemic bubbles can form with no ill intent, through ordinary processes of social selection and community formation. We seek to stay in touch with our friends,

who also tend to have similar political views. But when we also use those same social networks as sources of news, then we impose on ourselves a narrowed and self-reinforcing epistemic filter, which leaves out contrary views and illegitimately inflates our epistemic self-confidence.

To explain how an epistemic bubble could manifest itself in the practical world, we can imagine a person called Anna that belongs to the right-wing of the political spectrum. This is mainly due to her inheriting her parents' right-wing values through her conservative upbringing, and she has in turn chosen friends with more or less the same values as herself. The informational sources Anna is exposed to on a day-to-day may include whatever is featured on her news feed on Facebook (where she only has her friends and family as contacts). On Facebook, the news she will be exposed to will almost only come from right-wing sources, firstly because of her right-wing network, and secondly because of Facebook's own algorithms that fill her news feed with posts based on her interests. With the latter, Anna has not actively excluded informational sources writing about news from left-wing sources criticizing right-wing politics. Also, she would not necessarily discredit these sources if she should suddenly run into them. Nevertheless she, in virtue of her social selection, has never run into opinions challenging the right-wing, and she lacks relevant voices in the contemporary political discussion. This could also be the case if one made the example a bit less extreme, including sources such as news channels on TV, a newspaper subscription and even having a job with colleagues holding different opinions than her. There are news networks that run stories and opinion pieces in line with Anna's already accepted beliefs, and the same goes for newspapers, so she could follow both without getting any closer to bursting her bubble. Her colleagues might very well disagree with Anna's right-wing opinions, but knowing what Anna is interested in and which beliefs she holds, they could very well avoid any topics touching these beliefs as to sustain some collegial friendliness.

It seems intuitively problematic that an epistemic agent can find herself in such a position. Thi Nguyen identifies what this problematic trait is by pointing towards Goldberg's coverage-reliability (Thi Nguyen, 2020, 143), or more specifically, that epistemic bubbles lack it. Coverage-reliability is "[...] the completeness of relevant testimony from across one's whole epistemic community" (Thi Nguyen, 2020, 143). It is obvious that Anna does not have a completeness of relevant testimony, seeing as her epistemic community solely consists of people conveying information that is in line with right-wing politics and perspectives, as well as people who are not bothered going into political discussions with her at all. Epistemic

bubbles lack a representative conception of the available ideas, positions, opinions and facts that are available, and are thus in a quite poor position to land on well-formed beliefs.

After introducing coverage-reliability, Thi Nguyen narrows his definition of epistemic bubbles to "[...] a social epistemic structure which has inadequate coverage through a process of exclusion by omission" (Thi Nguyen, 2020, 143). He continues by stating that this omission is encouraged by two primary forces. "First, there is an epistemic agent's own tendency to seek like-minded sources" (Thi Nguyen, 2020, 143). In our example this is apparent through Anna befriending and networking with others that share her right-wing values and opinions. Having a network of sources solely speaking the mind of the right-wing provides, as we already have uncovered, a bad coverage-reliability. Thi Nguyen goes on to name the second force, namely that "[...] there are the processes by which an epistemic agent's informational landscape is modified by other agents" (Thi Nguyen, 2020, 144). This modification can take the form of censorship, state media control, algorithms etc. (Thi Nguyen, 2020, 144). In the case of Anna, any chance of her being exposed to sources other than those from the right-wing on Facebook would be through Facebook's own suggestions for new pages and persons to follow. The problem is that these suggestions are based on algorithms that have figured out what Annas interests are, namely right-wing politics, and therefore would feed her with the arguments and opinions that she already supports. The algorithms can be seen as a recommendation from a fellow partisan rather than an attempt on challenging her views or broadening her information spectrum.

Thi Nguyen notes that it is important to distinct epistemic bubbles from echo chambers (Thi Nguyen, 2020, 141-142), and by building on the work of Kathleen Hall Jameison and Frank Capella he defines echo chambers as the following: (Thi Nguyen, 2020, 142):

An echo chamber [...] is a social epistemic structure in which other relevant voices have been actively discredited. [...] [A]n echo chamber's members share beliefs which include reasons to distrust those outside the echo chamber. Echo chambers work by systematically isolating their members from all outside epistemic sources [...] By discrediting outsiders, echo chambers leave their members overly dependent on approved inside sources for information.

A contemporary example of an echo chamber is closed, neo-Nazi internet forums. They often begin with creating a narrative about the white, Christian and sometimes even European core family, conservative values and an "us vs. them" group mentality. Once one agrees with this narrative, one uses the group mentality to discredit sources that are not a part of the group,

naming them enemies of the cause, as well as amplifying the credibility of the insiders. By calling media houses and other voices for Zionists, Marxists, or multiculturalists, they become an enemy to the neo-Nazi cause, and therefore not to be trusted. In the end the only sources members of echo chambers of this kind are the neo-Nazi authorities themselves, as well as any blog, newspaper and other voices that agree with their narrative. The consequence is then that members of the echo chamber unconditionally rejects all counterarguments and relevant voices that are from outside the echo chamber itself, which clearly can be seen as an epistemic vice.

Now that we have established what epistemic bubbles and echo chambers are according to Thi Nguyen, we can move towards the question on whether they are possible to escape. Thi Nguyen writes that we also must ask ourselves "[...] whether there is an escape route that we might reasonably expect an epistemically virtuous agent to discover and enact" (Thi Nguyen, 2020, 144). The point of the question is that even if we theoretically can imagine an escape route, it is important to find out whether an agent is able to be aware of said escape route.

Exiting an epistemic bubble

To proceed we can first look at epistemic bubbles. Is it possible to escape an epistemic bubble, and can we reasonably think that an epistemic virtuous agent could and would do so? As previously noted, Thi Nguyen thinks epistemic bubbles are quite fragile. Given that an epistemic bubble consists of excluded information without said information being discredited, all that we need to shatter the bubble is exposure to it (Thi Nguyen, 2020, 154). Moreover, Thi Nguyen holds that an epistemically virtuous agent would search informational sources outside her usual network, and therefore the escape route is both discoverable and actable (Thi Nguyen, 2020, 154). If we recall right-wing-Anna, she could be content with only being exposed to news and information with a right-wing angle. If that is the case, it could possibly be correct to call the failure to search after other informational sources for an epistemic vice. However, if she is an epistemically virtuous agent, she would gather her news from other sources in addition to her Facebook network, and would then break out of her epistemic bubble.

Therefore, on the first hand, it seems like leaving an epistemic bubble is quite unproblematic. Exposure to other sources of relevant information is all it takes, and getting this exposure is as simple as broadening your search for information. But is it this simple in all cases of epistemic bubbles? Arguably not. Firstly, one could imagine someone who finds themselves within an epistemic bubble without having access to a wide variety of informational sources. If you only have real access to a single source about a topic, broadening your search for information about said topic will undoubtedly be difficult. Furthermore, it also seems wrong to claim that a person in this epistemic bubble is blameworthy for it. To be able to ascribe blame to a person for an action or an omission, the person should have been able to act differently in some way. This is not the case whenever a person is in an epistemic bubble due to lack of access to other relevant information. If you want to learn some basic astronomy and your *only* available source on the topic is Plato's *Timaeus*, you will undoubtedly end up in a faulty astronomical epistemic bubble, yet you cannot really be blamed for this. You did not have any other sources to consult, no other possible course of action, so there was no other possible result than ending up in that epistemic bubble. Secondly, there might be some cases where there are no other information or relevant sources about a certain topic. If epistemic bubbles function as omitting relevant information and sources, they seem to entail that there is something to omit. Therefore, in cases where there conceivably is no information to omit, one must either grant that one is not in fact in an epistemic bubble when one relies on a narrow (yet all-encompassing) set of sources or refine the definition of epistemic bubbles so that these cases are excluded.

It seems relevant to discuss whether either of these two cases constitute epistemic bubbles at all. Again, if epistemic bubbles are when you omit relevant information about a certain topic, omission must take place. Yet, regarding the first case, if you do not have access to other informational source about a certain topic and have already consulted those sources you *do* have access to, there really is not talk of any omission. You cannot omit that with you do not have access to. In such cases, however, one could possibly argue that one then should exhaust all options to *gain* access to new relevant information, and this argument seems to hold some merit. Yet, we can at the very least conclude that in cases where you 1) do not have access to other sources and 2) do not have any conceivable way to *gain* necessary access, then you have not omitted relevant information and, therefore, are not in an epistemic bubble. A similar conclusion seems to be sensible regarding the second case. To omit something, there must be

something to omit. It then follows that if there is nothing to omit, then an omission cannot occur. So, there is no epistemic bubble here at all.

Therefore, we can for now conclude that epistemic bubbles in fact are quite unproblematic to leave. One simply must expose oneself to the relevant information and sources that one previously has omitted. On a personal level there might of course be some difficulty to motivate this exposure, as one might be more than comfortable enough within one's bubble but given that one tries to leave an echo chamber the chances of success seem high. Further, the exceptions we discussed – where one does not have appropriate access to relevant information and where there is no other relevant information – arguably are not echo chambers at all, and therefore no hinderance for conceiving echo chambers and their exit strategy as Thi Nguyen suggests.

Exiting an echo chamber

In his discussion on the possibility of exiting an echo chamber, Thi Nguyen begins his argument by distinguishing between two types of echo chamber members. Firstly, you have an agent that willingly joins an echo chamber. Such an agent would be epistemically blameworthy according to Thi Nguyen, because "[f]or an agent in full possession of a wide range of informational sources, to abandon most of them and place their trust in an echo chamber [...] is surely some form of epistemic vice" (Thi Nguyen, 2020, 154). Since entering the echo chamber in this case is voluntary, one would most likely also be able to use the same power of choice to leave it again.

However, the view that willingly entering an echo chamber necessarily is an epistemic vice does not seem entirely convincing. To explain the objection here, I would first start to outline how a process of entering an echo chamber willingly could look like. Firstly, one already holds a set of beliefs about a topic, as well as trust in several relevant sources on the topic and/or distrust in several relevant sources. It is worth to note that this trust might very well be justified. Secondly, to protect yourself from the sources you do not trust, or possibly in an attempt to only rely on the sources you in fact do trust (the result being seemingly the same), you find social niches where your already-held beliefs less likely will be scrutinized.

Now, the first stage of willingly entering an echo chamber is not at all malignant and cannot necessarily be called an epistemic vice. A good epistemic agent will of course have a set of beliefs as well as different relevant sources one trusts and distrusts. There might of course be an epistemic fault in how that trust or mistrust is founded, but an epistemic agent has surely not committed any error so far given that the levels of trust are sufficiently justified. Thi Nguyen seems to place the epistemic vice with the second step, where one enters a group where your already-held beliefs do not come under scrutiny. One of the main points he makes here is that by removing your beliefs from possible revision or scrutiny from sources outside whichever group you chose will make your beliefs more prone to error and counteract any attempt to reach the truth. Yet, it does not seem entirely convincing that this necessarily is so.

One example that might counter Thi Nguyen's conception of the necessarily malignant echo chamber is that of Odysseus and the Sirens. Knowing that the Sirens' song would prove fatal to Odysseus and his crew, he ordered his crew to put beeswax in their ears so that they would not hear the song, whilst he was tied to the mast of his boat so that he could not yield their call. So, his crew actively excluded the Sirens' call given their trust in Odysseus' command, and thus spared their lives. This example shows us how sometimes, exclusion can be a necessary tool in order to protect oneself. In the case of Odysseus and his crew, the exclusion served to protect them from death, which is a very terrible consequence of merely listening to what someone has to say. The potential consequences do not have to be equally dire in order for protection by exclusion making sense. Some sources are manipulative, faulty, full of fallacies, or have some other quality that can be epistemically harmful. Digesting information that come from such sources may potentially contribute to damaging your prospects of gaining and retaining knowledge, substituting knowledge with poorly founded doubt, logically incoherent arguments, or just factually wrong beliefs. Would it not make good sense to exclude these sources from one's informational intake, at least not granting them the same consideration as other sources, if these sources have plausible potential to do epistemic damage? One might very well respond to this question by saying that, following an established free-speech view, we should invite all voices to the table no matter how faulty their opinions may be. If what they are saying is wrong, we will be able to get a better grasp of why we are right. And if they are right, we get the chance to correct our own mistakes. I will not argue that this view regarding how to deal with potentially epistemically harmful sources is wrong in all cases, but I do still hold that there are some cases where excluding such sources can be wise. Odysseus' case is one, but one could also imagine an authoritarian

regime which heavily utilizes propaganda in order to affect its citizens' opinions. If one is aware that the propaganda channels in fact are propaganda channels, I argue that one should follow Odysseus' example and stuff one's ears with beeswax.

Thi Nguyen does not consider our objection, and rather moves on to considering cases where he does not believe agents can be blamed for partaking in an echo chamber. Based on the example Thi Nguyen presents (Thi Nguyen, 2020, 154-155), we can imagine a child *S* that is raised by parents that are themselves in an echo chamber – in our example a fundamentalist Primitive Baptist church. *S* grows up learning from her parents, which a child reasonably could see as an epistemic authority (Thi Nguyen, 2020, 154). In such a case, Thi Nguyen argues that "[...] when that child eventually comes into contact with the larger epistemic world [...] the echo chamber's beliefs are fully in place, such that the teenager discredits all sources outside of their echo chamber" (Thi Nguyen, 2020, 154-155). When *S* goes out into the world after a childhood dominated by the teachings of her echo chamber, *S* will meet the world based on background information based on these teachings. Thi Nguyen points out that *S* could be a reasonable epistemic agent that seeks, investigates and evaluates new sources, but all new sources will be measured after how it fits with *S*'s background information. If the background information then in turn has taught *S* the opposite of what the outside sources say – as well as not to trust outside sources - then the result will be that *S* gets a confirmation regarding the outside sources' untrustworthiness as well as a strengthening of her own theories brought down from the echo chamber (Thi Nguyen, 2020, 155). If *S* grew up believing that the earth is between six and ten thousand years old, as well as believing that the mainstream media and scientists are untrustworthy, she will find it problematic to meet the outside world where most claim that the earth is billions of years old. Using her background information regarding the source and her own belief about the earth's age, she will probably end up reconfirming that outside sources merely are heathen blobs. This usage of background information to judge new sources is, as previously mentioned, quite rational given *S*'s upbringing. "The worry here is that agents raised within an echo chamber are, to through no fault of their own, epistemically trapped" (Thi Nguyen, 2020, 155).

This self-confirming belief process *S* and other agents that have grown up in echo chambers are subjects for makes the question on whether it is possible for themselves to see that they are a part of an echo chamber seem negative. However, Thi Nguyen proposes one possible escape route out of such echo chambers. This route is grounded in Thomas Kelly's discussion

of belief polarization. "Belief polarization is the tendency of individuals, once they believe that p , to increase their belief that p " (Thi Nguyen, 2020, 156). Then, if we again think about the aforementioned S that believes that p the earth is between six and ten thousand years, this will mean that once S starts believing p as a child she will from then on will only strengthen her belief that p . One of the reasons behind this tendency is that once someone has acquired a belief, they tend to be more critical to counter-arguments to said belief, and at the same time overlook flaws with supporting arguments (Thi Nguyen, 2020, 156). So, whenever our S is approached with arguments claiming that the world in fact is billions of years old, she will tend to try to scrutinize and find flaws with those arguments. And whenever someone that holds and argues for the same belief p that she does, she will merely let it pass. "Thus, their critical reflection is likely to reinforce previously held beliefs" (Thi Nguyen, 2020, 156). This, according to Kelly per Thi Nguyen, violates the *Commutativity of Evidence Principle*: (Thi Nguyen, 2020, 156)

[T]o the extent that what it is reasonable for one to believe depends on one's total evidence, historical facts about the order in which that evidence is acquired make no difference to what it is reasonable for one to believe.

The violation made by belief polarization is clear. If S first encounters the belief that p , and she ends up believing that p , she will have a tendency to just strengthen her belief that p despite her being met with a whole lot of counter-arguments to p later in life. When S keeps believing that the world is eight to ten thousand years old despite the overwhelming amount of counter-arguments she meets after the belief is formed, the order one receives evidence in matters (Thi Nguyen, 2020, 156).

Although belief polarization makes the chance of getting out of an echo chamber look steep, Thi Nguyen argues that there is an exit to be found in the Commutativity of Evidence Principle. If one wants to escape from the echo chamber, one has "[...] to undo the influence of the historical ordering of their encounters with the evidence" (Thi Nguyen, 2020, 157). To do this, one would have to throw out all of their background knowledge and restart the entire knowledge-gathering process (Thi Nguyen, 2020, 157). Our aforementioned S would have to forget her beliefs regarding the earth's age, as well as all other background knowledge, and start from scratch. Of course, actually doing this would have its practical difficulties. Thi Nguyen points out that is very similar to Descartes' method, namely epistemically starting over again, not trusting any sources and requiring full certainty from beliefs (Thi Nguyen,

2020, 157). Thi Nguyen calls this the *Cartesian epistemic reboot* and goes on to saying that it is not feasible given that it is more or less impossible without trusting external testimony (Thi Nguyen, 2020, 157). Imagine if our *S* started over again, now not believing in any her background information (including that the earth is some thousands of years old). On her quest on gaining new beliefs, she would not get far at all without the testimony of others. Given this problem with the Cartesian method, Thi Nguyen attempts finding an escape route through something he calls the *social epistemic reboot* (Thi Nguyen, 2020, 157). We can better understand what such a reboot is by starting to identify what Descartes' project was based on, namely his concern about all the falsehoods he had grown up believing and the later beliefs based on these falsehoods (Thi Nguyen, 2020, 157). The problem in a case like Descartes' is that what Thi Nguyen calls *credentialing structure* one has from one's upbringing is flawed (Thi Nguyen, 2020, 157). This flawed structure will in turn influence new beliefs which will be flawed as well, and in a way be self-confirming in that it strengthens the credentialing structure (Thi Nguyen, 2020, 157). In this way, one cannot help someone out from an echo chamber by targeting one belief at the time, as this structured web of beliefs is, in a way, codependent and the beliefs both support and rely on each other. Thi Nguyen therefore still holds that the subject trapped in the echo chambers has to epistemically start over, just not in the way Descartes did (Thi Nguyen, 2020, 157).

So how would one go forth to conduct a social epistemic reboot to eliminate the influence from the historical order one has received evidence? Well, as with Descartes, the agent has to start by temporarily set aside *all* beliefs, especially those originating from one's credentialing structure (Thi Nguyen, 2020, 157). Now that the agent has blank slates, she does not have to dismiss the testimony of others, nor does she have to live up to Descartes' demand for absolute certainty. However, Thi Nguyen holds that she has to restart socially, that is, not let their previous credentialing structure get away without scrutiny (Thi Nguyen, 2020, 157). Explained through our previously mentioned *S*, she will have to set aside all her beliefs from her upbringing all those beliefs that in turn have sprung out of it. When this is done, she has to meet the world with all its arguments, opinions and beliefs unbiased and without laying any weight on what she previously held as beliefs. That means that the beliefs of the church community she originally lived cannot escape critique, while sources like mainstream media, scientists and encyclopedias no longer are immediately discarded. Thi Nguyen compares the state the agent has to be in to the one of a cognitive newborn, a state "[...] of tentative, but defeasible, trust in all apparent testimonial sources" (Thi Nguyen, 2020, 157-158). The result

of meeting the world in this state would, most likely, be for our *S* to believe that the earth in fact is round and lead her out of the echo chamber she was born into.

The social epistemic reboot seems a bit easier to accomplish than the Cartesian one, but not by far. Nevertheless, Thi Nguyen holds that it is not completely unrealistic, a claim he supports by pointing to the story of Derek Black who was born into a neo-Nazi family, but still ended up abandoning the belief system he grew up with (Thi Nguyen, 2020, 158). Supposed that Thi Nguyen's account of this exit strategy is correct, he still has not answered whether one could reasonably expect an agent to reboot or listed any reasons for why an agent even would want to. His objection starts with the example of Black, who ended up befriending a Jewish co-student who gained Black's trust. It was through his friend he ended up realizing that many of his background beliefs were false, which successively made him leave the whole platter (Thi Nguyen, 2020, 158). Thi Nguyens point is that "[s]ince echo chambers work by building distrust towards outside members, then the route to unmaking them should involve cultivating trust between echo chamber members and outsiders" (Thi Nguyen, 2020, 158). The key to helping a person out of an echo chamber does not lie in approaching them with neutral facts because, as we already have seen, echo chambers work through discrediting external sources. Instead, one must try to overrule the discrediting made by the echo chambers, gaining the trust of their members and then drag them out of it (Thi Nguyen, 2020, 159). Meeting our prior *S* with documents from the field of astronomy to persuade her in believing that the earth is round will not work since it opposes her credentialing structure. Before giving her the facts, we will have to overcome the fact that she does not see us as a credible source and try to gain her trust.

It here seems appropriate to separate between the cognitive and affective levels of our beliefs. Generally, but maybe especially in echo chambers, we do not adopt or hold on to beliefs solely because of our cognition. The motivation one may have to enter or remain in an echo chamber does not be solely of cognitive value. In addition to our cognition, affective values seem to have a role to play. We may have an emotional attachment to our beliefs which makes it harder to scrutinize them, requiring an excessive number of counterpoints to dislodge it. Further, we may have an emotionally loaded judgement of different sources, making us trust some or distrust others based on largely affective criteria rather than cognitive ones.

Epistemic vices

In this section, I will introduce the notion of epistemic vices as described in Quassim Cassam's "Vices of the Mind" before I consider whether a member of an echo chamber can be said to have epistemic vices.

Epistemic vices are, according to Cassam, "[...] personal intellectual failings that have a negative impact on our intellectual conduct" (Cassam, 2019, 2). I may have good reason to believe that I could not jump over a three-meter-long gap between two solid rocks, but overconfidence could challenge that belief and leave me with bruises. Here, overconfidence serves as such an intellectual failing, which stood in the way of me acting on a justified belief regarding my jumping abilities. My overconfidence is therefore an epistemic vice. It is important to note that although epistemic vices are obstacles to knowledge, that does not entail that all obstacles to knowledge are epistemic vices. (Cassam, 2019, 3) There are several factors that can weigh in on our epistemic endeavors that are not epistemic vices, for instance medical and psychiatric symptoms etc. Psychosis, insomnia, and other such conditions can very well be an obstacle to knowledge but can still not be said to be an epistemic vice. This is because such conditions do not merit criticism, which is an important component to epistemic vices according to Cassam. (Cassam, 2019, 4)

In order to reach a deeper understanding of what epistemic vices are, we will start by considering the blameworthiness of epistemic vices. It is first worth noting that Cassam's account focuses on epistemic blame rather than moral blame, epistemic blame not necessarily leading to moral blame. (Cassam, 2019, 123) Cassam's starting point is the claim that one generally cannot be blamed for what we are not responsible for. (Cassam, 2019, 123) The responsibility spoken of here is not one of answerability, that is, that we are able to cite reasons for why we have certain characteristics or act in a certain way, but rather one of accountability, that is, that we can be held accountable for our characteristics or actions. (Cassam, 2019, 124). What, then, does it mean to be accountable?

Cassam suggests two possible answers, namely that accountability here should be understood as either having 1) acquisition responsibility or 2) revision responsibility for our epistemic vices. (Cassam, 2019, 124) Acquisition responsibility has to do with whenever an agent's actions and omissions have led to her developing an epistemic vice, the thought being that she

then is responsible for the vice coming into being. (Cassam, 2019, 124). Revision responsibility has to do with cases where an agent potentially can revise or counteract her vice. (Cassam, 2019, 124) This notion of responsibility stems from the thought that it does not make sense assigning blame to those who could not possibly have acted differently, as we have previously discussed in our treatment of Thi Nguyen's paper on echo chambers and epistemic bubbles. This also opens for the possibility of assigning blame to those who cannot be said to be blameworthy for *acquiring* an epistemic vice, but who nevertheless conceivably are able to *revise* said vice. A type of agent that then can be rightfully assigned blame is someone who inherits epistemic vices from their upbringing to no fault of their own, but that nevertheless comes into a position where they are able to amend or correct their vice, only to then failing to do so. Thus, coming from a racist household or community does not necessarily make you blame-free if you are able to revise your racially motivated prejudice as you get older.

So, we can be said to be blameworthy for an epistemic vice if we have acquisition responsibility or revision responsibility for said vice. Further, we have revision responsibility for a vice if we are able to control it. Let us take a closer look at what revision responsibility entails by asking what it entails to have control. Cassam suggests three different answers, namely: (Cassam, 2019, 124-126)

1. Voluntary control
2. Evaluative control
3. Managerial control

Voluntary control is when you could do something right away if you chose to do so, and according to Cassam we lack this type of control over our own beliefs (Cassam, 2019, 125). Evaluative control, in the realm of belief, has to do with our capability to amend our beliefs in accordance with our own evaluations. (Cassam, 2019, 125-126) If I believe that it will take seven minutes to boil a soft-boiled egg, but after following my own belief end up with a hard-boiled egg, I should be able to adjust my belief in accordance with my new evaluation based on my experience of boiling eggs. In that way, I have exerted evaluative control, thus altering my own belief. Lastly, managerial control is “when we manipulate some ordinary object to accord with our thoughts about it” (Cassam, 2019, 126), something which according to Cassam is possible also when it comes to our own beliefs (Cassam, 2019, 126). Managerial

control over our own beliefs can take the form of looking for evidence to support a belief I wish to hold (confirmation bias), or different belief management tactics such as hypnosis or “positive thinking”. (Cassam, 2019, 126)

Applying the different types of responsibility and control to a single belief seems straightforward. For instance, take an agent that holds the false belief that “The earth is flat”. We now know that if she is to be blamed for holding the false belief, she must be responsible for it. She can be acquisition responsible for the belief if she started believing it by her own actions. She could for instance limit her informational input about the shape of the earth to flat-earth sources, and over time start to shift her belief because of repeated and exclusive exposure to flat-earth information in combination with effective conspiratorial and manipulative rhetoric. If she 1) believed the earth was round prior to her engagement with the flat-earth sources, 2) was aware of the flat-earth communities having such an effective way of communicating their beliefs and 3) that such communication could have the effect of convincing people of holding beliefs that they prior to the exposure held to be false, it seems correct to call the agent acquisition responsible for her new false belief. She could also be revision responsible for the belief that the earth is flat, insofar she has the level of control necessary to revise or change her belief. Voluntary control is, as already noted, difficult when we speak of beliefs. If an agent believes that the earth is flat, you cannot simply wish the belief away. You might be motivated to change your beliefs for any number of reasons, but we do not have direct voluntary control over our beliefs. We may however have evaluative control over a given belief. If you believe the earth is flat, but then are met with information stating that the earth in fact is round, you may very well revise your initial belief based on a new evaluation of the matter. You do not necessarily need some new empirical evidence in order to motivate this change of belief either, as you also could re-evaluate your belief that the earth is flat and notice that you were faulty in your original evaluation and therefore should no longer believe that the earth is flat. The flat-earthier also seem to have, in principle, managerial control over her belief, insofar that belief is not immune to hypnosis, therapy or other belief management tactics.

So, assigning types of control to different cases of holding a single belief is quite straightforward. Yet, despite how straightforward deciding on assigning responsibility to an agent for holding a single false belief is, false beliefs may be the result of a more systemic flaw in the form of an epistemic vice. If, for example, an agent comes to believe that the earth is flat

after hearing it from a single source with no authority on the subject, one might say that the adoption of this belief stems from an epistemic vice of gullibility. One could say that the agent should have consulted more authoritative sources on the topic, or at the very least consulted more than a single source for this claim. Does this agent have responsibility for her gullibility? It generally does not seem plausible to say that someone is acquisition responsible for one's gullibility, so unless an agent willingly and knowingly increases their own gullibility (insofar that is possible), gullibility is rarely comes about by one's own actions. The agent could possibly be revision responsible for her gullibility if she is able to control it. Does she have either voluntary, evaluative, or managerial control over her gullibility? She does not seem to have voluntary control over her gullibility, since one cannot simply stop being gullible at will. She may however, at least in some sense, have evaluative control over her gullibility. Given a certain level of cognitive status, an agent that comes to know that she previously has landed at the wrong conclusion due to her gullibility could evaluate the situation as "my gullibility led to that mistake". This is not really control in the same sense as when we consider evaluative control in relation to a single belief. It is possible to revise a single belief using evaluative control, as in the case with the boiled egg. In the case of epistemic vices, it seems that evaluation does not really lead to one being able to revise one's vice, but rather just that one is able to acknowledge that the epistemic vice is there. The gullible person that time after time is led astray may evaluate those cases and find out that she had been too gullible, and in turn be motivated to become less motivated. Yet, this evaluation alone does not alone lead to her being able to control her gullibility. She can come to acknowledge her gullibility, but then she has just ended up at the same position as the one who wishes to voluntarily control their vice: an acknowledgement of what needs to be changed but still without the possibility to just *will* the change. More powerful tools have to be used in order to actually initiate a change away from the vice, and maybe managerial control can be this tool.

Now, does our agent have managerial control over her gullibility? Arguably she does. If she recognized that she in fact is able to (a recognition that for instance can come from her previous evaluations), and in addition to that is motivated to become less gullible, then presumably she could work on strategies to that end. (Cassam, 2019, 132) As previously stated, managerial control involves a form of belief management. Possible strategies involving such belief management could for instance be practicing being more skeptical, not taking anything at face value but instead weighing up different opinions on any given subject.

(Cassam, 2019, 132) Even if practicing healthy skepticism may be incredibly hard for the gullible agent, external help such as family, friends or a therapist could help with the burden. (Cassam, 2019, 132)

It does, however, not seem entirely correct to assign managerial control to a gullible agent that does not know that she is gullible. She would of course still be able to practice skepticism, but if she does not know that she is gullible then why would she? Cassam writes that if an agent has managerial control over her vices, this implies that she not only must be able to manipulate that vice but also be aware that she has the vice at all, as well as being motivated to change. (Cassam, 2019, 129) As previously stated, evaluating previous experiences where your vice somehow led you into trouble, for instance gullibility leading you to wrong beliefs, may make you aware that you in fact are gullible. If the trouble is serious enough, it might even motivate you to do something about the gullibility. (Cassam, 2019, 129) Another source of motivation might be your own negative evaluation about the vice. If you look down on gullibility and those who are gullible, you might be motivated to work towards avoiding being so yourself. (Cassam, 2019, 129) How do we then consider managerial control and, in turn, blameworthiness and responsibility in cases where the agent lacks self-awareness and/or motivation to change? Cassam writes that even if an agent possibly could work their way away from a vice, they do not in fact have managerial control if they are oblivious to them having said vice. (Cassam, 2019, 129) So, the possibility of being ignorant to one's own epistemic vices opens for a situation where an agent has no voluntary, evaluative or managerial control over their vice. Can, then, one omit an ignorant agent for both blame and responsibility? According to Cassam, this depends on *why* the agent is ignorant to her own vices. (Cassam, 2019, 130) He continues by outlining three different reasons for such ignorance, reasons I choose to call:

1. Culpable self-ignorance
2. Self-ignorance by lack of motivation
3. Full self-ignorance.

Firstly, culpable self-ignorance is when an agent both has access to evidence that supports that she in fact has a vice and understands said evidence, but still does not believe that she has a given vice since having it would counter how she regards herself. (Cassam, 2019, 130) Here, the agent is ignorant because of her own self-deception, and can thereby possibly be both

responsible *and* blameworthy for her vice. Secondly, self-ignorance by lack of motivation is when an agent in fact does know she has a vice but lacks motivation to do something about it. (Cassam, 2019, 120) Now, agents who are self-ignorant by lack of motivation may be blameworthy according to Cassam, but only if the lack of motivation is due to some factor(s) that the agent is to blame for. Cassam points specifically to lack of motivation “[...] due to depression or some other condition for which she is blameless” (Cassam, 2019, 130) as vindicating. Thirdly, there is full self-ignorance, something which we will treat in a sub-chapter of its own.

Stealthy vices

How about cases where an epistemic agent truly is ignorant about their vices, that is, full self-ignorance? Here, Cassam introduces the notion of “[e]pistemic vices that obstruct their detection [which] are inherently hard to detect” (Cassam, 2019, 145), a type of vice he calls *stealthy vices*. (Cassam, 2019, 145) A central idea behind Cassam’s account of stealthy vices is the *Dunning-Kruger effect*:

[...] in many social and intellectual domains incompetent individuals lack the skills needed to know that they are incompetent. As a result, they have inflated views of their performance and abilities (Cassam 2019, 144)

The Dunning-Kruger effect shows what full self-ignorance could look like, namely “[b]eing too incompetent to understand one’s own incompetence [...]”. (Cassam, 2019, 144) This does not entail that one is too incompetent to self-scrutinize by attempting to reflect over one’s own possible epistemic vices, but rather that one would not actually learn anything about one’s stealthy vices from this reflection, and that the reason you would not learn anything is *because of* said stealthy vices. (Cassam, 2019, 145) An agent who possesses one or more stealthy vices can reflect around themselves and their character all they would like without discovering their flaws, they are simply incapable to unveil them. And, importantly, this incapability is a result of the very flaws in question. Cassam continues the accounts on stealthy vices by discussing several questions. We will, in the next, focus our attention on two of them, namely: (Cassam, 2019, 147)

1. What makes an epistemic vice stealthy? What are the obstacles to knowledge of one’s own epistemic vices, and which are the stealthier or stealthiest vices?

2. How, if at all, is it possible to overcome the obstacles to knowledge of our stealthy vices?

Obstacles to vice detection

As we already have seen, what separates a stealthy vice from other epistemic vices is that they are difficult to detect. Cassam suggests that the difficulty in detecting stealthy vices generally can be caused by four different obstacles: (Cassam, 2019, 148)

1. Laziness or fear
2. Lack of the necessary conceptual resources
3. Vice contradicting self-conception
4. Vice-denial.

Firstly, laziness or fear as an obstacle for vices detection simply has to do with an agent not being bothered or being too afraid to reflect over their possible vices. (Cassam, 2019, 148) An inflated self-image and arrogance could possibly be a reason for not bothering, whilst fearing what one might find once one starts an introspective vice-analysis could be the reason for not doing it out of fear.

Secondly, an obstacle to vice detection can be “[...] that one lacks the conceptual resources necessary to know one’s own epistemic vices” (Cassam, 2019, 148). Simply put, if an agent is not aware of the concept of their vice, for example what dogmatism actually is, one presumably cannot know that one *is* dogmatic either. It is generally hard to pin something down when one does not have the language or conceptual understanding to describe it, and this is true also when it comes to vices.

Thirdly, an obstacle to vice detection can be that the detection of a given vice would contradict one’s self-conception and self-image. (Cassam, 2019, 148) The example Cassam uses to illustrate this is an agent having a full grasp of the concept of dogmatism, while simultaneously dogmatically believing herself to be undogmatic. (Cassam, 2019, 148) The agent so rigidly believes that she is a certain way that it overshadows any signs of the opposite being true.

Fourthly, vice detection becomes difficult if an agent does not regard the vice in question as a vice at all. (Cassam, 2019, 148) Cassam uses an example of a racist in a racist community where racism is not remotely conceived as a negative trait. (Cassam, 2019, 148) If the tenants of a vice are deemed as acceptable, or even applaudable, by most people around you, it becomes very difficult to be able to acknowledge that your epistemic behavior actually is vicious.

Overcoming the obstacles to vice detection

How could one break through all these possible barriers to detecting one's own epistemic vices? Could engaging in *active critical reflection*, that is, “[r]eflection on one's own epistemic vices [...]” (Cassam, 2019, 148), be a possible strategy? If I am unknowingly close-minded, could I not find out about my own closed-mindedness by reflecting upon my vices? According to Cassam, when it comes to stealthy vices, the answer is no. That is since the very vices one would unveil if the reflection was a success *counters* the virtues necessary for its success. (Cassam, 2019, 149) If I am closed-minded and try to reflect over my own vices, I would need to be open-minded in my reflection. Unfortunately, however, I am close-minded, and thus my closed-mindedness blocks my chances for finding out that I am closed-minded. (Cassam, 2019, 149) Here we see an important quality of stealthy vices, namely that they are inherently hard to detect. (Cassam, 2019, 149)

So, if active critical reflection is out the window as a possible vice detection strategy, what about if one simply is told that one has a certain vice? Could we rely on testimony as a way in which we could come to know which vices we possess? Cassam holds that this does not work either, since we ultimately may rely on active critical reflection even when we gain knowledge from testimony. (Cassam, 2019, 155) If someone tells me that I am closed-minded, my closed-mindedness might affect my active critical reflection so that I do not fully consider the testimony about my closed-mindedness. If my closed-mindedness makes me dismiss views that counter my already accepted beliefs, then hearing someone call me closed-minded does not necessarily get me any closer to realizing that I am closed-minded. The same goes with epistemic overconfidence, for instance. If I am epistemically overconfident and someone tells me that I am epistemically overconfident, my overconfidence can affect my active critical reflection about that testimony. What active critical reflection may tell me is that I am not epistemically overconfident, I am just justifiably confident in my own epistemic

behavior, and the person telling me that I am epistemically overconfident is plainly wrong, maybe because she is not as epistemically superior as myself. Stealthy vices are hard to crack, again since they block their own detection. Hearing that one has a stealthy vice from someone else does not change the fact that it is the hearer that must reflect over the statement, and that the stealthy vice is ever present in this reflection.

So, we have seen that stealthy vices are inherently hard to detect, and that neither active critical reflection nor testimonial knowledge about our vices seem to a viable route towards detection. Yet, according to Cassam, we do not yet need to abandon our pursuit for an answer on how we may come to know our stealthy vices. (Cassam, 2019, 157)

The viable route suggested by Cassam starts off by returning to the Dunning-Kruger effect: that one can be too incompetent to register one's own incompetence. Cassam cites a study that showed that increasing one's competence led to them being able to acknowledge that one in fact *had* been incompetent *prior* to their increase of competence. (Cassam, 2019, 157) Does this mean that an increase in competence can lead to understanding one's own shortcomings, including one's stealthy vices? Not quite, since what is gained by the increase in competence is understanding how one *used* to be, not how one in fact is here and now. (Cassam, 2019, 158) If we think of stealthy vices, it seems quite possible for the formerly closed-minded or the formerly epistemically overconfident to understand that she *used* to be closed-minded or epistemically overconfident. Yet, that is not really the challenge we are discussing. The challenge is making the person that is *currently* closed-minded or epistemically overconfident understand that she *is* closed-minded or epistemically overconfident. What is needed is what Cassam calls here-and-now vice knowledge. (Cassam, 2019, 158)

How does one achieve the here-and-now knowledge necessary to overcome the obstacles to vice detection, then? Cassam suggest an interesting cause of here-and-now knowledge, namely traumatic experiences. He defines traumatic experiences as the following: "A traumatic experience is typically 'a sudden, unexpected, and potentially painful event' that 'ruptures part of our way of being or deeply held understanding of the world'" (Cassam, 2019, 159). In this rupturing way, any knowledge gained from traumatic experiences are gained accidentally, as opposed to any knowledge gained from active critical reflection. (Cassam, 2019, 159) Traumatic experiences can be breakthrough experiences in that they are potentially able to break through one's self-ignorance and thus revealing a stealthy vice. (Cassam, 2019,

159) When one experiences something as a shock, it can give them sudden realizations about who and how they are, realizations that would otherwise be hard to come by. Moreover, these realizations would not necessarily be accessible through active critical reflection, or at least not active critical reflection alone. (Cassam, 2019, 159) Cassam uses the example of an unknowingly gullible person going bankrupt because of his gullibility, where the shock is so big that he manages to realize that he in fact is gullible. (Cassam, 2019, 159) Had it not been for the bankruptcy he might very well have stayed unknowingly gullible for the rest of his life, but the shock and trauma following from the bankruptcy was so significant that he managed to break through his own self-ignorance. What is important to note about traumatic experiences is that they *can* lead to a break-through, but that they do not necessarily do so. (Cassam, 2019, 159) The main point is not that traumatic experiences guarantee vice-detection, just that they sometimes can.

Cassam describes traumatic experiences that serve as breakthrough experiences through three different notions: (Cassam, 2019, 160)

1. As transformative experiences
2. As quantum change
3. As cataleptic impressions.

We will now take a closer look at the first and third notions, as they relate the most to our discussion on echo chambers.

Traumatic experiences can be transformative experiences in that they can give you an insight that one could not have gotten elsewhere, or at least would not have without that experience. (Cassam, 2019, 160) One can make a distinction between personally and epistemically transformative experiences, where personally transformative experiences changes who you are while epistemically transformative experiences change your perception of who you are. (Cassam 2019, 160-161) A personally transformative experience can, for instance, be having a child. (Cassam, 2019, 160) Having a child can change your outlook on life, your priorities, and your values. The formerly mentioned unknowingly gullible person can be said to have had an epistemically transformative experience when he went bankrupt. If the bankruptcy led him to understanding that he in fact is gullible, the transformation lies in a shift in his self-

perception. His beliefs about himself go through a change from before and after the experience, and this change is set in motion by the experience. (Cassam, 2019, 159)

Nussbaum via Cassam defines cataleptic impressions as impressions which “[...] by their own experienced quality, certify their own veracity” (Cassam, 2019, 162). Further, “[...] by assenting to a cataleptic impression we get ourselves into ‘a condition of certainty and confidence from which nothing can dislodge us’” (Cassam, 2019, 162) Nussbaum uses the example of Proust’s Marcel not understanding that he is in love with Albertine until he hears of her departure, making it a realization that he did not reach by using a regular reasoning process but rather through emotion. (Cassam, 2019, 162-163) This is an important quality of the traumatic experiences Cassam discusses and supports the claim that traumatic experiences can provide insight that active critical reflection cannot. (Cassam, 2019, 163)

There is more to be said regarding traumatic experiences in the context of stealthy vices, such as discussing how traumatic experiences can lead to lasting insight of one’s stealthy vices and, hopefully, lasting change. Yet, knowing the basics about the promise of traumatic experiences as a cause of vice-detection is sufficient in the context of this thesis. In the next, we will move on to looking at Cassam’s discussion on self-improvement for those knowingly guilty of epistemic vices.

Prospects of self-improvement

Now, say that one comes to know that one has a stealthy vice, for instance through a traumatic experience as described in the previous. What happens then if one does not want to settle for just *knowing* that one has a vice, but also wants to do something about it?

Cassam makes a distinction between two ways of answering this question: *optimism* and *pessimism*. The optimistic view is that self-improvement in relation to epistemic vices both is possible and often accessible. (Cassam, 2019, 169) Still, important prerequisites for self-improvement are knowledge of having the epistemic vice, motivation for self-improvement and possessing the necessary means for self-improvement. (Cassam, 2019, 169) The pessimistic view doubts that epistemic vices are malleable, or at the very least malleable in that one can make lasting change to one’s epistemic vices, and instead suggests an entirely different strategy to dealing with the problem of epistemic vices. (Cassam, 2019, 170) Instead

of self-improvement, instead of doing something with the epistemic vices themselves, they suggest *outsmarting* the vices: “[...] that is, finding ways to work around them so as to reduce their ill effects” (Cassam, 2019, 170). One example of that kind of out-smarting tactic is workplaces that receive resumés where one has anonymized the applicants’ name and gender in order to outsmart potential prejudiced hiring policies, for instance favoring a specific gender or being biased against foreign applicants.

Now, Cassam’s own account of epistemic vices does not follow either of the routes but is rather compatible with both, he claims. (Cassam, 2019, 170) “My claim isn’t that lasting self-improvement is always possible but that it sometimes is” (Cassam, 2019, 171) He holds that although we sometimes will have to do with just outsmarting an epistemic vice, that does not make all attempts at countering and reducing epistemic vices hopeless. And this hope may lie in cognitive therapy.

Cassam suggests there are several kinds of cognitive therapy that can be a passable route towards self-improvement. (Cassam, 2019, 178) He uses an example based on himself to show how this is so. Cassam wrongly predicted the results for the 2016 presidential election and Brexit referendum, which lead him to conclude that his predictions were based on wishful thinking. (Cassam, 2019, 178) He was then both aware of his vice of wishful thinking and motivated to do something about it, and in future similar situations he asked himself more skeptical questions and questioned his own predictions in a larger degree than before. (Cassam, 2019, 178) Of course this does not necessarily guarantee that he will make predictions less based on wishful thinking in the future, but engaging in this self-scrutiny as an attempt to combat his vice *might* contribute to self-improvement.

There are of course several factors that can seriously complicate any attempt at self-improvement. One’s epistemic vices could ultimately be the result of the environment around you, thus making self-improvement difficult if not impossible for the individual until something is done about her surrounding environment. (Cassam, 2019, 182) Also, there is a case to be made for claiming that self-improvement is a much harder task for those who probably need that self-improvement the most. (Cassam, 2019, 183) Working on one’s epistemic vices can be an immensely difficult endeavor. After all, it involves restructuring your own rational self, often tearing up incredibly deeply rooted negative traits. But despite how difficult it is, the important thing to take away from this is that self-improvement *is*

possible. It might not be possible for all, but for many it can be a walkable path. (Cassam 2019, 187) And when the potential gains from such self-improvement are of such significance, this is very good news for anyone who values an epistemically virtuous approach to knowledge.

Epistemic vices and echo chambers

In this section, I will try to explain echo chambers considering Cassam's account on vices and, more specifically, stealthy vices. The questions we will be considering in the next is whether there are any vices to talk of when considering echo chamber members, and if so, are they stealthy vices? We will also be taking a new look on the example of an echo chamber member's successful exit in order to better understand both epistemic vices, echo chambers and necessary preconditions for an exit.

Epistemically vicious echo chambers

If we recall the chapter on Thi Nguyen's work on epistemic bubbles and echo chambers, an echo chamber is a structure where relevant voices that are external to the echo chambers are *actively discredited*, thereby isolating its members from external informational sources. Some sources are deemed as untrustworthy, malignant, or manipulative, and the echo chamber members therefore refrain from considering views or statements from these sources at all. Are there any vices in action in this discreditation process? I argue that there often are. In the next I will describe how an echo chamber's epistemic mechanisms might ultimately be a reflection of different underlying epistemic vices.

Firstly, members of echo chambers can potentially be closed-minded. Members of an echo chamber's active exclusion of different sources of information is a lack of willingness to consider opinions and information countering their already held beliefs, and they are thus closed-minded. And as long as the echo chamber which the members belong to does not happen to have gotten everything right (all their beliefs being correct while external information could not possibly bring in any new knowledge), this closed-mindedness undoubtedly serves as an obstruction for knowledge.

Secondly, I will make the argument that echo chamber members can be guilty of dogmatism. We will define dogmatism as “[...] the uncritical, partial, and possibly irrational persistence of some opinion” (Dogmatism, Blackwell, 2009). We can regard the echo chamber member as uncritical in virtue of their already defined closed-mindedness. If you do not even consult sources with information that challenges the held beliefs of the echo chambers (given that such sources exist and one has feasible access to these), your belief in the echo chamber’s established beliefs is arguably uncritical. Through the active discreditation of external sources, the unwillingness to consult those outside the echo chamber itself, one denies oneself any opportunity to using one's critical faculties. Further, the echo chamber member is also partial in their dependence on chamber-internal informational sources and their active discreditation of chamber-external sources.

Thirdly, prejudice may also be part of echo chamber members’ set of vices. Endre Begby defines prejudice as “[...] a negatively charged stereotype, targeting some group of people, and, derivatively, the individuals who compromise this group” (Begby, 2021, 8-9). Of course, prejudice can be a component of the echo chamber in the sense that they collectively hold one or more prejudiced beliefs. The prejudiced beliefs that are most central in our discussion on echo chambers are those who contribute to an epistemic community being labelled as an echo chamber at all, namely prejudiced beliefs that serve as justification for excluding and discrediting certain sources or groups of sources. If an echo chamber excludes and discredits outsiders on the basis of the outsiders belonging to a group which the echo chamber hold negatively charged stereotypes against, and that the exclusion and discreditation is *because of* the content of that stereotype, that echo chamber becomes guilty of prejudice.

Fourthly, I argue, members of echo chambers have inflated epistemic confidence. Again, this stems from them actively discrediting external sources. The motivation for this discreditation can differ from chamber to chamber. Some may do it partly or solely due to a protective standpoint, as in that they want to protect themselves from potentially damaging voices (as with Odysseus and the sirens). In this sense, one cannot necessarily hold that the epistemic confidence is a case of epistemic overconfidence, given that one’s confidence is justified. This seems to be the case for Odysseus, who makes his crew stuff their ears full of beeswax to protect their ability to make rational decisions. Others may actively discredit (and thus show their inflated epistemic confidence) due to faulty or poorly formed beliefs about the outsiders, for instance due to prejudiced beliefs about those excluded.

Thus, we see that members of an echo chambers arguably have several different epistemic vices, and that all these vices can be connected to how echo chambers actively discredit chamber-external sources. We can therefore see that the difficulty of dealing with echo chambers largely lies with this hugely important component of actively discrediting informational sources, and that this discreditation both can stem from as well as maintain a large set of epistemic vices. It therefore seems appropriate to continue our discussion by taking a closer look at the vices in questions, and to take separate looks at vices in accordance with whether they are stealthy or non-stealthy.

Understanding Thi Nguyen's proposed exit strategy through the vice-lens

Now, clearly members of echo chambers can be guilty of several epistemic vices. To get a clearer understanding of which role epistemic vices can play in echo chambers, we will now re-examine Thi Nguyen's proposed exit strategy while bearing the account of epistemic vices in mind. Recall, in writing about Derek Black, a former neo-Nazi who understood that his convictions were wrong after befriending a Jewish co-student, Thi Nguyen concludes that an exit strategy for echo chamber members can be to cultivate trust between echo chamber members and those outside the echo chamber. (Thi Nguyen, 2020, 148) In the following we will see how the example of Black can be understood in terms of vice theory, how this new understanding challenges Thi Nguyen's proposed exit strategy, and possible amendments to said exit strategy.

We can assume that Black previously was guilty of at least prejudice, epistemic overconfidence, dogmatism, and closed-mindedness. In their thoughts on racial supremacy, neo-Nazis have an ideological prejudice against those who do not belong to their categorized "supreme" race. Not only is this ethically reprehensible, but it is also a clear obstacle for the acquisition of knowledge. The epistemic consequence of such prejudice includes missing out on any possible knowledge conveyed by those who do belong to the so-called inferior races. This is because for the neo-Nazi, those belonging to the so-called inferior races cannot be trusted by virtue of their race or religious background. This distrust may also be placed with anyone who has been perceived as supporting the so-called inferior race, or those who are seen as manipulated by them.

We can understand parts the worldview of the neo-Nazi as having a distinction between the Us and the Other, the Other being those belonging to or advocating for the so-called inferior races (or even just not actively supporting “the Us”). It seems natural then to say that neo-Nazis’ prejudiced view on the Other (and resulting inflated view of the Us) consequentially leads them to being guilty of a form of epistemic overconfidence. They do not necessarily think that the Us is epistemically infallible (although they can think that), but rather that the only ones who can be trusted are part of the Us. This judgement nevertheless inflates the epistemic confidence in those belonging to the Us since it ultimately leads to the belief that one generally cannot trust, and thus not learn anything from, the Other. This serves as a huge obstacle to knowledge, making much otherwise obtainable knowledge unobtainable through the exclusion that follows one’s epistemic overconfidence.

The prejudice, epistemic overconfidence, and dogmatism typical in neo-Nazi echo chambers (and other epistemically vicious echo chambers) are all intertwined with closed-mindedness. It is not always easy to say something about the causal connection between the different vices. Prejudice against the Other can lead to closed-mindedness, in that the consequence of this prejudice generally leads one to be less likely to consider the views of the prejudiced group(s) as well as anyone regarded as in support of or cooperation with the prejudiced. Likewise, closed-mindedness can lead to prejudice, if one’s closed-mindedness leaves one over-exposed to negative stereotypes of a given group and without being open for countering views. Epistemic overconfidence can be linked with closed-mindedness in that epistemic overconfidence makes you less prone to consult or consider information and opinions that counter your already held beliefs, since the epistemic overconfident agent is so sure that she already possesses the correct and relevant information. Likewise, if you are closed-minded and shy away from countering opinions, your trust in your already-held beliefs is likely to strengthen due to the lack of scrutiny which in turn may make you epistemically overconfident. Further, closed-mindedness can lead to a form of epistemic overconfidence that lasts even if one stops being closed-minded. We saw earlier how an epistemic agent S broke the Commutativity of Evidence Principle. Through closed-mindedness she adopted a faulty belief that the world is eight to ten thousand years old without getting any counterevidence. As time goes this belief may get stronger due to the lack of counterevidence, potentially making S epistemically overconfident about her belief. If she then, later on, stops being closed-minded, and exposes herself to loads of counterevidence against her belief, her

epistemic overconfidence can combine with a sort of confirmation bias so that she does not change her belief.

We can conclude, then, that Black and other members of epistemically vicious echo chambers are guilty of several epistemic vices. Now if, as Thi Nguyen claims, exiting an echo chamber requires building trust with those outside the echo chamber, the question becomes how someone guilty of the array of epistemic vices listed above build trust with outsiders? If one is closed-minded, one will normally not come into situations where even starting to build such trust is possible. This is due to the previously discussed stealth of closed-mindedness. Willingly entering situations or conversations that possibly could be the start of a trust-building process is by definition unlikely for the close-minded person, and this *because of her* closed-mindedness. Echo chamber members could of course, and often do, end up in situations where they find themselves talking to someone who disagrees with the echo chamber's adopted beliefs, but close-mindedness can interfere here as well by making the echo chamber member disregard any countering views presented, being too closed-minded to fully consider those views. Epistemic overconfidence might also fuel this lack of consideration of contrasting views.

So, being guilty of epistemic vices makes it hard for an echo chamber member to break out. If building trust with outsiders is necessary for breaking out of an echo chamber, then being guilty of epistemic vices serves as a terrible starting point for any breakout attempt. Yet Black, who was guilty of several epistemic vices, managed to break out of the echo chamber by befriending a man of Jewish faith. How could this be? Thi Nguyen does not discuss this in depth himself but understanding echo chamber-membership in terms of epistemic vices might enable us to better understand how Black broke out of the echo chamber. We cannot know for certain how Black's process was, which epistemic vices he had or the strength of his epistemic vices, so we will in the next hypothesize from this example. What we *do* know is that his newly founded friendship with a person belonging to a group he had firmly held stereotypes against played a large part in his chamber-exit.

Firstly, he was placed in a situation where he had to interact with an outsider. Sharing a dormitory with a Jew was not an active choice for Black and would by all probabilities never be an active choice he would make himself given his closed-mindedness and prejudiced beliefs targeting Jews. As we have already discussed, closed-mindedness is a stealthy vice, a

vice that prevents its own detection. Because an agent is closed-minded, she will avoid situations etc. that has the potential to show that she is closed-minded. That is true insofar the closed-minded has a choice in the matter, but in Black's case that choice was probably made over his head. Someone else made the decision on who he should share a room with. Taking the choice out of Black's hands could then contribute to bypassing his closed-mindedness, at least a closed-mindedness so strong that he would not wish to share a space with a member of the so-called inferior groups.

Secondly, in conversations with outsiders, one does not only have to speak about conversation topics which the echo chamber has clear opinions on. We can assume that one of those clear opinions is, in Black's case, of course that his dormitory roommate is necessarily inferior following his religious background but given that this belief and his general closed-mindedness apparently was not strong enough to leave college once he learned about his roommate's religion, their co-living opens up for communication between the two. This communication could include talking about topics that Black's echo chamber does have clear opinions on, but this would serve as a bad foundation in order to build trust. Yet, there are loads of conversation topics available that does not revolve around those echo chamber beliefs. Insofar these "permitted" conversation topics, that is, those which does not make an echo chamber's closed-mindedness or epistemic overconfidence kick in as a kind of blocking mechanism, are sufficient in order to create a sort of personal relationship between the echo chamber member and the outsider, there does not seem to be any conclusive reason for why this personal relationship could not be accompanied by trust between the two.

If one successfully manages to build trust through engagement in these permitted conversation topics, then the web of Black's interconnected epistemic vices just might fall. He might eventually see that the prejudice he held against Jews does not hold up when it comes to his roommate. Experiencing close-handedly that his prejudiced beliefs did not hold up in one case might also hurt his *dogmatic* trust in those beliefs, leading him to the realization that if there is one exception to how he believes Jews are, there might be other exceptions out there. In turn, this realization is a sort of epistemic failure in that he understood that he was at least somewhat wrong in his beliefs regarding Jews. This failure can, firstly, serve as a kick in the teeth, deflating his epistemic confidence. It shows that he is not epistemically infallible, and that there in fact is at least one belief he at one point held to be completely true which later revealed itself as false. Secondly, the realization that Black is not epistemically infallible

might just serve as a motivation to become less closed-minded (although that motivation does not necessarily follow from epistemic failure). If one has been sure that one is completely right about something just to later realize that one in fact was wrong about it and has a desire to believe that which is true, one might be motivated to search for countering views and arguments outside one's preexisting epistemic community. This paragraph describes at least somewhat how the process was for Black in his departure from his neo-Nazi echo chamber, understood in terms of epistemic vices. And the snowball that started rolling first was the overriding of his closed-mindedness that happened when he was placed in the same dormitory as his new friend.

As we have already seen, Cassam suggests that traumatic experiences can be way to detect one's stealthy vices. This is because traumatic experiences can make one get here-and-now vice knowledge in a way active critical reflection cannot. (Cassam 2019, 158) Surely, Black did get a dosage of here-and-now vice knowledge when he figured out that he had a bunch of false beliefs from his echo chamber, but did this knowledge stem from a traumatic experience? When traumatic experiences are understood as "[...] a sudden, unexpected, and potentially painful event [...]" (Cassam, 2019, 159), the answer seems to be "maybe". Black's experience of moving in with and then befriending his Jewish roommate can probably be said to be unexpected, but the building of the trust necessary to form a friendship is hardly a sudden endeavor, at least not in the sense bankruptcy or physical assault is sudden. Also, the experience itself was not necessarily painful (although the realization of believing in a set of false beliefs might have been). Yet, there are some connections between Cassam's account of traumatic experiences and what we imagine happened in Black's case. Firstly, we can understand Black's experience as a transformative experience. In befriending his roommate, he got an insight that he would not have gotten without that experience. Chances are that he would not have suddenly stopped believing in his Neo-Nazi beliefs without that newly founded friendship, or that he would have realized that his beliefs were in fact wrong in another way. It was the friendship that made him understand that he was in the wrong, and he would not have understood that without the friendship. Secondly, we can understand Black's experience as being a cataleptic experience. Quite like Marcel fully realized that he was in love with Albertine when he heard of her departure, a realization that he could not be dislodged from, Black seems to have realized that he was in the wrong about at least the characteristics of Jews when he befriended his roommate. And again, this experience is what set in motion the dismantlement of his web of Neo-Nazi belief.

Are all echo chambers malignant?

We have already established that echo chambers include active discreditation of relevant voices. This discreditation is often closely tied to a variety of epistemic vices, such as closed-mindedness or epistemic overconfidence. However, it may also be possible that echo chambers may exist where such active discreditation does not necessarily relate to any epistemic vices. In other words, it is possible that there may be echo chambers where the members do not exhibit any of the epistemic vices we typically think are a part of echo chambers. If this is the case, then it may be necessary to reevaluate and perhaps even revise the current definition of echo chambers as proposed by Thi Nguyen. In this section, I will attempt to make a case for the existence of such non-vicious echo chambers and explore the implications of this potential revision.

Academics as an echo chamber

Let us start out by considering an example of how members of a university institute may, epistemically speaking, behave in their work. These employees are specialized in one or more fields and need to stay updated on new information, the latest research and academic articles on these fields in order to do their jobs. Yet, in their search for new publications on their field of expertise, they are likely to pay more attention and consideration to peer-reviewed academic articles than an opinion piece in the local newspaper. Full consideration could be granted to an opinion piece given that it is written by someone from within the research community, someone who has ascended a similar academic ladder to the institute employees themselves. The same goes for extending full consideration to verbal presentations, where one is more likely to consider what is said in a lecture by a professor than what is said by an unqualified public speaker on the same topic.

The institute employees seem to discredit external sources, more specifically sources that are not peer-reviewed articles or publications and lectures by authors that has not been peer-reviewed in other contexts. It is maybe not a discreditation in the way that they refuse to listen external sources, but rather that they do not fully consider them as they would consider an academic peer. This, I argue, is in line with Thi Nguyen's definition of echo chambers. Can

we then, drawing on this example, say that the institute employees in fact are members of an echo chamber?

As per Thi Nguyen's definition of echo chambers, if we are to say that our institute employees belong to an echo chamber we will have to be able to call the above-mentioned exclusion and discreditation active, and those who are discredited will have to somehow be relevant voices. Further, the institute employees must share a belief that there are reasons to distrust those outside academic circles. In addition to that, the institute employees must be systematically isolated from non-academic sources. Lastly, resulting from the listed criteria, the institute employees must end up overly dependent on the approved academic sources. (Thi Nguyen, 2020, 142) Are all these criteria present in the example of our institute employees?

First, let us discuss whether the described exclusion and discreditation in fact is active. The activity here is to be understood as being the opposite to inactive discreditation methods such as omission, the latter being how epistemic bubbles work. While omission is when one happens to leave out sources, unintendedly, for instance because your Facebook-algorithm does not think you would be interested in a set of sources (Thi Nguyen, 2020, 143), active exclusion is when a source is "[...] actively assigned some epistemic demerit, such as unreliability, epistemic maliciousness, or dishonesty" (Thi Nguyen, 2020, 146). One discredits a source based on these epistemic demerits and grant them less or no consideration in order to maximize one's chances to gain true knowledge. This seems to be a plausible description of how our institute employees engage with sources. In focusing on peer-reviewed journals in one's search for information, our institute employees arguably attempt to find information that is correct, reliable and trustworthy. The general idea is that peer-reviewing serves as a control-mechanism to check whether a paper holds a certain level of quality. The implication here is that when a paper has not been peer-reviewed one simply does not know whether the paper is reliable, and thus should be read with a bigger portion of skepticism than if the paper had been peer-reviewed. To put it in other terms, peer-reviewed journals are seen as a reliable source, while non-peer-reviewed articles and papers are seen as having a larger potential for unreliability. Due to this difference in trust placed in peer-reviewed and non-peer-reviewed papers, and rooted in a general goal of gaining knowledge, non-peer-reviewed articles are discredited, and this discreditation is active.

Secondly, let us consider whether the discredited voices are relevant. For this to be true, some of the voices that are being excluded by our institute employees will have to have something to say about their field in question that is not covered by those who do make it through the peer-reviewing process (thus limiting the institute employees' coverage-reliability). This may in turn lead to the field missing out on new takes and new arguments on their subject matter. (Thi Nguyen, 2020, 143) Are there, then, any relevant voices being left out by our institute employees? Conceptually, for this to be true, all it takes are someone out there who has something new to say about the institute employees' subject matter that is not granted consideration by the academic community. The truth-value of whatever those someone has to say is beside the point, because as long as 1) the opinion is practically accessible by the institute employees, 2) the opinion potentially could contribute to increasing the institute employees' coverage-reliability or 3) the opinion could provide just one single perspective or information of value to the field, it seems that the opinion is by definition relevant. I will simply assume that the coverage-reliability of the academics is not infallible, and hold that at least some fields within the academics are discrediting at least some relevant voices.

Thirdly, let us consider whether there in fact is a shared set of beliefs for our institute employees that state reasons for why non-academic sources cannot be trusted. I will not dwell on this point for too long, but I argue that those beliefs are in place, and that this follows from the first point regarding the exclusion that takes place through the difference trust placed in peer-reviewed and non-peer-reviewed papers. The shared belief is that peer-reviewed articles generally are more trustworthy than non-peer-reviewed articles, and thus non-peer-reviewed articles should be granted less trust.

Fourthly, let us consider in which fashion our institute employees are isolated from non-academic sources, and whether this isolation can be regarded as systemic. This fourth point also relates to the aforementioned difference in trust. I argue that our institute employees indeed are isolated in the sense described by Thi Nguyen, because of their preference of academic sources over non-academic sources. This preference is, again, explained by our discussion of the trust given to academic and non-academic sources. It must be said, however, that, the isolation is not a full isolation, and of course our institute employees can both access and extend full consideration to non-academic sources. But this is of course also the case for echo chambers such as Neo-Nazi communities, they too can both access and extend consideration to echo chamber-external sources. This isolation is not, I argue, a impenetrable

isolation that no outside sources can overcome. Rather, it has to do with members of an epistemic community generally avoiding outside sources for the reasons considered in the discussion regarding why outsiders cannot be trusted. One could still engage with outside sources, but then with an increased sense of skepticism. Thus, I argue that this isolation is more of an insulation, an insulation that is supposed to serve as protection against sources generally deemed as untrustworthy.

To finish our examination of whether our institute employees can be said to be in an echo chamber, we will briefly consider the factuality of the previous points lead to our institute employees being overly dependent on academic sources. The sense of being “overly” dependent should not, I argue, be read as necessarily serving as an obstacle to knowledge. Thi Nguyen does not himself go into much depth about what this overly dependence entails, but I suggest understanding it as generally utilizing and considering the echo chamber-internal sources while simultaneously being, as Jamieson and Cappella suggest via Thi Nguyen, “[...] highly resistant to any outside sources” (Thi Nguyen, 2020, 146). Further, I argue, this dependence entails that it is *because of* this resistance to outside sources that one generally utilizes and considers echo chamber-internal sources. When assuming this conception of overly dependence, one could in fact maximize one’s chances of gaining knowledge even though one is highly dependent on a set of sources while being resistant to outside sources, granted one gains more true knowledge through this dependence than one would without it. *Odyssey and the Sirens* is an example of this. So, on the basis of the previously discussed points, are our institute employees overly dependent on academic sources? I argue yes. This does not damage the prospects of holding that the institute employees in fact are maximizing their chances of gaining knowledge through engaging in active discreditation and exclusion of sources. Instead, as we have already seen, this simply shows that our institute employees generally utilize and consider academic sources, and that this tendency is caused by the resistance they have to non-academic sources. In turn, this resistance is caused by the institute employees’ shared belief that academic sources generally are more trustworthy than non-academic sources, combined with the more general goal of gaining knowledge. The shared belief might very well be true, and if it holds up then the dependence on academic sources might be epistemically sound. Yet, given that it in fact is true, this is perfectly compatible with being overly dependent according to our assumed definition.

Oddly enough, one might say then, it seems that there is a case to be made for claiming that our institute employees in fact are members of an echo chamber. Thi Nguyen's definition of echo chambers can then, arguably, fit both university institutes as well as Neo-Nazi communities. Yet, intuitively, it seems quite wrong to call the epistemic structures of the two the same. In the next I will shed light on an important distinction that separates the two, enabling us to get a better grasp of the importance of talking about epistemic vices in the context of echo chambers.

Academics as seen through the vice-lens

Even though both the university institute and the Neo-Nazi community may exhibit the characteristics of an echo chamber as defined by Thi Nguyen, it is important to note that the epistemic structures of these two entities are not equivalent. In other words, while both may actively exclude and discredit relevant sources, the underlying justification for this is vastly different. Here, I will describe how different the two are by utilizing Cassam's account on epistemic vices. We can start off by considering which, if any, epistemic vices are the institute employees guilty of.

Firstly, one could argue that they are guilty of a sort of closed-mindedness, in that they are unwilling to grant full consideration to academic-external sources on their subject of expertise, and usually look solely to peer-reviewed articles and the works of other persons in the field to gain new information on the subject. The philosophy professor that wishes to update herself on the contemporary debates around Hegel's thinking will be more likely to check out a couple of academic journals rather than visiting the discussion forum Reddit's Hegel-forum, where anyone with an internet connection can publish content. The scope of sources which the institute employees fully engage with is limited to academic-internal forums and publications, while non-academic sources are disregarded or granted less consideration.

The closed-mindedness described here stems from the institute employees' level of epistemic confidence. The consequences of their levels of epistemic confidence are like the ones of echo chamber members such as Black, namely thinking that one knows better than those external to one's epistemic community and thus granting less consideration to the outsiders' views. The institute employees grant more epistemic value to academic sources than non-academic

source since they hold that academic sources are more dependent, trustworthy, methodologically in line with one's own way of thinking of a certain subject, and therefore ultimately more likely to be able to transmit new truthful knowledge. For the institute employees, those who belong to academic institutes have a higher epistemic standing than non-academic sources and limiting one's scope to consulting and considering academic sources regarding academic subjects is generally considered as sufficient in one's search for information. To sum up, there is a case to make for academics having a high degree of epistemic confidence in the academic community. Yet, it seems odd to name the epistemic confidence of the institute employees as guilty of *overconfidence*. To illustrate the oddity, we can compare the justification for the epistemic confidence shown by Neo-Nazis and the epistemic confidence shown by our institute employees. As previously discussed, Neo-Nazis can be guilty of epistemic overconfidence as a consequence of how they discredit sources that are not sympathetic to Neo-Nazism or ideologies and central ideas close to Neo-Nazism, such as white supremacy, anti-Semitism and Islamophobia. In other words, their discreditation is closely connected to another epistemic vice, namely prejudice. The discreditation of external sources is fueled by a belief that such external sources are untrustworthy, which makes way for the implicit belief that true beliefs are generally (or at least more likely) to be found within the Neo-Nazi circles themselves. This inflates the epistemic confidence levels of those inside those circles and serves as an obstacle to knowledge. Members of academic circles, our institute employees for instance, generally also have a high degree of epistemic confidence in the circles they belong to. Yet, the justification for this confidence is vastly different to the one of Neo-Nazism. In the next, I will argue that the difference in justification not only shows that the epistemic behavior of the institute employees in fact does not constitute closed-mindedness and epistemic overconfidence as Cassam describes the two terms, but also that this in turn suggests a new definition of echo chambers.

First, let us consider *why* our institute employees have such high epistemic confidence. This *why* is also the answer to why they do not extend the same level of trust and consideration to non-academic sources as they do academic sources. One possible answer is quite simple, namely that they generally in fact do know better, and generally are the best sources available. This is because of most academic fields' methodological approaches over time being refined to maximize its ability to uncover truthful statements within its own field and new information being peer-reviewed to ensure its quality. The institute employees are, in different ways, trained to become more epistemically virtuous through their academic training, and

peer-reviewing is an example of verifying that the epistemic virtues that are praised in academic circles are followed. Therefore, academic sources are extended more trust and consideration than non-academic sources, mainly because one has better reason to expect that they are following good scientific practices, practices grounded in a variety of epistemic virtues, which in turn maximizes the chance of those sources leading the reader closer to knowledge. One such epistemic virtue is arguably a sort of discernment one can practice because of one's deep knowledge of a field. People trained within the academics might also generally be expected to be more objective in their engagement with the subject matter than non-academics. Further, the academically trained can be expected to have an overview over the state of the art of their subject and knowledge of their fields' canon, an expectation that might not be equally strong when it comes to non-academics. Given this knowledge, one might then expect academics to be better suited to place a piece of knowledge in a larger context, show how it relates to other pieces of knowledge and thus perhaps bring about a better understanding of whatever the topic is. There might be other virtues one could expect from our institute employees as well, but the main point is that it seems plausible to say that we have higher epistemic expectations and demands to academics than non-academics, in virtue of research institutes and other academic institutions' fundamental goal of gaining knowledge. They do not necessarily always meet these expectations and demands, failure is always a possibility, but they are generally expected to perform better and more epistemically virtuously than non-academics. At the very least, academics themselves tend to hold this belief. Given that they in fact are epistemically virtuous and generally are well-versed in their subjects, academics seem to have good justification for their high epistemic confidence. And if they are aware of where their competence ends, not believing that they know more than they in fact do, it does not seem to be right to call this a case of epistemic overconfidence.

The institute employees' justification for their high epistemic confidence levels may also provide us with justification for their apparent closed-mindedness. If one reasonably can believe that the best, most well-informed, objective and epistemically virtuous opinions on a certain subject is to be found within the academics, the epistemically virtuous agent would naturally focus their search for knowledge on the subject here. If we assume, as we have done, that the academics in fact have a justified belief in their ability to convey information in an epistemically virtuous way, and that one generally has better reason to trust the academics' ability than the ability of non-academics, it could very well be the epistemically sound choice to practice closed-mindedness in the sense we have discussed. Further, it must be said, the

institute employees' previously described closed-mindedness is not *truly* closed-minded, at least not necessarily. Public debate, critical papers and public discourse may convey a variety of opinions and pieces of knowledge, and academic circles may (and often do) both consider, trust, and engage with non-academic sources. Of course, the institute employees might prefer using academic sources over non-academic sources, but that does not entail that non-academic sources are not considered at all. They might, for instance, place great weight in a non-academic source when there are no academic sources on a given topic. Public non-academic debate might move faster than the academic debate, and in such cases, it makes perfect sense to grant full consideration to the non-academic voices. Granted, the academic circles might process and contextualize the debate in a different matter than the general public, but that does not take anything away from the fact that consideration is granted outside one's own academic epistemic circle. There are of course many other ways in which those belonging to the academics arguably engage with non-academics, but the main point is that there does not seem to be anything about the how our institute employees epistemically engage with different sources that entails that they would not open their minds for any pieces of knowledge granted that it contributes something to their field. They might prefer using academic sources, they might almost exclusively consult academic sources, but they do not necessarily put a hard ban on considering external sources.

Moreover, I argue, institute employees are not necessarily closed-minded at all, not even in a justified sense. Of course, individual members of a university institute may in varying degree be guilty of closed-mindedness, but they are not necessarily closed-minded, and those individuals who are closed-minded are not necessarily so because of their affiliation to the academics. Whole institutes and departments can also be guilty of closed-mindedness, for example if an institute has a goal on becoming a leading academic force within a certain scientific theory and pursue that goal so rigidly, so that they thereby end up missing out on potentially important knowledge stemming from viable but competing theories. Yet, the possibility of whole institutes and departments being closed-minded does not entail that they necessarily are so. The main point regarding both individuals and larger communities and the possibility of closed-mindedness is not that they cannot be closed-minded, but rather that exclusion of non-academic sources does not entail closed-mindedness. As we have seen, this exclusion is justified by the view that academics generally can be expected to convey epistemically virtuous and true knowledge, only making it natural to generally limit one's scope to academic sources. Recall, an important tenant of epistemic vices is that they serve as

an obstacle to knowledge. If the academic exclusion is done in order to maximize one's chances to gain knowledge, and that the exclusion in fact does so, it serves as an assistance rather than an obstacle to knowledge. Thus, the academic exclusion cannot be said to be an epistemic vice.

So, to summarize, I argue that our institute employees are not guilty of epistemic overconfidence, but rather have a justified high level of confidence. This level of confidence is justified because one can have a greater expectation of epistemic virtuous engagement with a subject matter from the academics due to their academic training. Further, the institute employees are not closed-minded either, at least not in the sense where closed-mindedness serves as an obstacle to knowledge. Of course, some members of the academics may be either or both closed-minded and epistemically overconfident, and the same goes for whole academic institutions, but this does not necessarily follow from academic practice. In other words, academic practice as we have described it does not necessarily entail epistemically vicious behavior.

Consequences of the existence of non-malignant echo chambers

So far in this chapter, we have been able to make two substantial conclusions. First, we saw that there is a case to be made for claiming that our institute employees partake in an echo chamber. Second, we saw that they partake in that echo chamber without necessarily being guilty of any epistemic vices. Thus, we can so far conclude that following Thi Nguyen's definition of echo chambers, there is room for calling epistemic structures that are not epistemically malicious echo chambers. In other words, there is such a thing as a non-malignant echo chamber.

Throughout this chapter, we have been examining the concept of echo chambers and their relationship to epistemic vices. By examining the specific case of the university institute and its members, we have been able to draw two substantial conclusions. Firstly, we established that there is a case to be made for claiming that our institute employees operate within an echo chamber, as defined by Thi Nguyen. Secondly, we saw that the members of the university institute partake in that echo chamber without necessarily being guilty of any epistemic vices.

In other words, we have motivated the possibility of the existence of echo chambers that are not caused by epistemically vicious behavior.

This leads us to the conclusion that, according to Thi Nguyen's definition of echo chambers, there is room for the existence of non-vicious echo chambers. In other words, it is possible for an echo chamber to exist without being inherently epistemically harmful. This conclusion challenges the conventional way in which one understands echo chambers.

A revised definition of echo chambers

Following from this insight, I suggest creating a distinction between malignant and non-malignant echo chambers. In discussions on epistemic structures such as echo chambers within epistemology, and perhaps within other fields, it will be appropriate to separate the two for the sake of a more pointed and accurate discussion. Maybe most importantly, I hold that it would serve the discussion well to have a clear conception of what the malignant echo chambers are, a conception that does not include non-malignant echo chambers. A focus on the malignant echo chambers also seems to be in line with how I believe the general public debate regarding echo chambers already is. Discussing the structures and exclusion strategies of an epistemic community might serve as some good intellectual fun but is arguably of little societal value in our context. Further, it seems intuitive to say that if we wish to discuss and find effective measures against epistemically malicious behavior, then it would benefit us to have a clear conception of how this malicious behavior functions.

So, instead of defining echo chambers as Thi Nguyen has done, where there is room for non-malignant echo chambers side by side with malignant echo chambers, I propose a rephrasing of the term *echo chamber* in order to better comply with public debate, ensuring that the discussion is focused on epistemically harmful structures, and separating those epistemically harmful structures from epistemic structures that are non-malignant, if not even epistemically virtuous. My proposed definition builds on Thi Nguyen's, but places much value on epistemic vices: "Echo chambers are structures where relevant voices that are external to the echo chambers are actively discredited, and this discreditation is justified in one or more epistemic vices".

Following this definition, the Neo-Nazi community Black came from remains an echo chamber. As we have already seen, Neo-Nazis discredit relevant voices on the basis of their ideological prejudiced beliefs regarding race, religion, culture etc. The Neo-Nazi epistemic community in question engages in a practice of exclusion and discreditation that is justified in prejudiced beliefs, and thus it is an echo chamber. The members of the university institute we have discussed also exclude and discredit outsiders. Yet, the justification of this process is based on the valuation of the tools and methods of academic voices, and the mostly accepted view that academic method is a good tool in order to gain knowledge. If academic sources generally are more versed in these methods, and one cannot expect those methods being used to the fullest extent outside the academics, it can make sense to generally focus one's attention to academic sources. One could of course always question whether this is not a prejudiced belief as well, and those within the academics would be wise to continue questioning this themselves, but except for those following the most skeptical lines of reasoning it seems safe to say that Neo-Nazi communities justify their exclusion in at least more prejudicial beliefs than members of an university institute do.

This revised definition of echo chambers could give us a clear method of identifying whether an epistemic community that discredits relevant voices is vicious or not. What is important in this identification is trying to understand whether the foundation for the community's practice of exclusion itself is based on epistemic vices or not. If it is, it is safe to say that the community in fact is an echo chamber, and one can then consider which vices are active in that echo chamber in order to figure out how one should deal with it. If it is not, the community hardly is an echo chamber. But even if it is not an echo chamber, it would be wise to remain alert about how that epistemic community handles its own sources, outside sources and how the discreditation process is justified, being aware of potential shifts that make the community so epistemically enclosed that it in fact becomes an echo chamber, thus potentially missing out on crucial information from relevant voices.

As we have seen in our treatment of Cassam's theory of epistemic vices, we have seen both that 1) although it may be difficult, detection of vices (including stealthy vices) is possible and 2) although it may be difficult, self-improvement *is* possible. This is a theoretical framework for ridding oneself of epistemic vices - if you are vicious: detect that you are vicious and get rid of the vice. With a definition of echo chambers that grounds the justification for the echo chamber's exclusion and discreditation process in one or more

epistemic vices, we can transfer Cassam's self-improvement framework to the realm of echo chambers.

When we discussed right-wing-Anna, we concluded that she was in an epistemic bubble, not an echo chamber. Because her exclusion was based on unintended omission rather than active discreditation and exclusion, she could not be said to be a part of an echo chamber following Thi Nguyen's definition. That fact does not change with our revised definition. That is not to say that Anna cannot be guilty of one or more epistemic vices. She could be regarded as closed-minded, or at least as epistemically incurious, insofar she has practically achievable access to non-right-wing sources but does not actively search for them. Right-wing sources having a monopoly on her news-focus could also lead to her being epistemically overconfident, this overconfidence coming from the lack of exposure to alternative views. The point about our revised definition of echo chambers is not that one cannot be epistemically vicious outside an echo chamber, but that one must be guilty of epistemically vicious behavior in order to be part of an echo chamber. Although epistemically vicious behavior that is not combined with an epistemically vicious form of active exclusion and discreditation still may be (and often is) blameworthy, that does not entail that one is part of an echo chamber. Our revised definition still requires the active exclusion and discreditation of relevant voices, and we have merely added a focus on *why* this exclusion and discreditation takes place.

If we again consider the example of Neo-Nazi communities, we can easily see how they are echo chambers under our revised definition as well. Neo-Nazi communities believe in a narrative of the Us versus the Other. The Us is usually white, Christian and conservative, while the Other are most people who do not belong to the Us, and especially Jews, Muslims and Marxists. An important feature of this narrative is the idea that the Other wants to somehow destroy or work against the Us and their cultural heritage, and that the Other possess some sort of inherent evil qualities. On the basis of this narrative, Neo-Nazis actively discredit and exclude relevant voices when those voices are a part of the Other. Moreover, these voices do not necessarily need to belong to a group the Neo-Nazis have clear negative opinions about. As long as one voices an opinion that somehow supports the groups the Neo-Nazis have a specified negative impression of, or an opinion that somehow criticizes the Neo-Nazis, one becomes a member of the Other by default. This distinction between the Us and the Other is what constitutes the Neo-Nazi's criteria of exclusion: those belonging to the Us

can generally be trusted and considered, while those belonging to the Other cannot. Of course, this distinction is based on false beliefs. Their ideas of Jews and Muslims and the rest of those belonging to the Other come from prejudiced ideas. As we have already seen, prejudice is an epistemic vice. And as our revised definition of echo chambers states, when a social epistemic structure actively discredits relevant voices on the basis of one or more epistemic vices, that epistemic structure is an echo chamber. Therefore, since the Neo-Nazi community 1) actively discredits relevant voices and 2) justifies this discreditation in prejudiced beliefs (and with prejudice constituting an epistemic vice), the Neo-Nazi community is in fact an echo chamber. Moreover, we understand *why* it is an echo chamber, and can target the epistemic vices that enable the echo chamber to exist. In our example we have stated prejudice as this fundamental epistemic vice, but there could of course be other vices functioning as justification for their engagement in discreditation and exclusion. By identifying all the fundamental vices, we can better understand why the Neo-Nazi echo chamber exists, and we can start making a more specific plan for how to combat it.

If we reconsider the example of the university institute, our revised definition shows why it does *not* constitute an echo chamber. I have argued that the way in which the members of the university institute interact with sources sometimes can be referred to as an active discreditation and exclusion of outsiders. This is due to how less consideration are given to sources that are not peer-reviewed or otherwise academically accepted. The university institute grants much more consideration to a peer-reviewed paper than an opinion piece in the local newspaper even if they discuss the same subject, or they consider the what the professor says in a lecture over what a public speaker says in a panel discussion. Yet, even if the members of the university institute do engage in an active discreditation and exclusion of relevant non-academic voices, their reasons for doing so are not necessarily epistemically vicious. Assuming that they exclude and discredit because of their justified faith in academic tools and methods, and the expectation that academic sources are more well-versed in using those tools and methods than non-academic sources are, this exclusion cannot be said to be justified in one or more epistemic vices. And since their justification for excluding and discrediting is not epistemically vicious, we cannot call the university institute an echo chamber.

So, in conclusion, our revised definition of echo chambers labels the Neo-Nazi community as an echo chamber, while leaving the university institute out of the definitions scope.

Conclusion

This thesis had as its aim to explore the concept of echo chambers in order to find out two things: what echo chambers are and what enables them to exist.

The first thing we did was to introduce and grasp the definition of echo chambers as presented by Thi Nguyen. His definition aimed first at making a distinction between two things that are often, and mistakenly, taken to be the same thing, namely epistemic bubbles and echo chambers. As we saw earlier, epistemic bubbles exclude relevant voices by omission and are quite easy to get rid of. In principle, all it takes to burst an epistemic bubble is broadening one's informational input by consulting more sources. An echo chamber is a much tougher nut to crack. This because they are structures where external sources are actively discredited. When you are in an echo chamber it does not really matter how broadly you find your sources, because you will not grant full consideration to a lot of those sources no matter what they say. Because of you and your echo chamber's shared belief about the untrustworthiness or unreliability of the echo chamber-external sources, you do not grant them consideration even when you are exposed to them.

Further, we quickly examined Thi Nguyen's own proposed exit strategy. Drawing on the example of the former Neo-Nazi Derek Black who ended up leaving his echo chamber after befriending his Jewish college roommate, Thi Nguyen suggested that what is necessary for an echo chamber exit is building trust with outsiders. It is because of the trust built between Black and his roommate that he was able to realize that he in fact was part of an echo chamber, and this realization enabled him to leave it.

After finding a preliminary definition of echo chambers and an understanding of how to leave them, we moved on to discussing Cassam's account on epistemic vices. Epistemic vices are "[...] personal intellectual failings that have a negative impact on our intellectual conduct" (Cassam, 2019, 2). Three further important notes on epistemic vices are that 1) they have to serve as an obstacle to knowledge (Cassam, 2019, 3), 2) you have to somehow be blameworthy if it is to be called an epistemic vice (Cassam, 2019, 4) and 3) you do not have to be responsible for the vice in order to be blameworthy for having it (Cassam, 2019, 123).

Although it can be a nuisance, detecting and fixing many vices is quite possible. Yet, there is a category of epistemic vices that are much more difficult to get rid of than the others, namely stealthy vices. The difficulty lies in that the stealthy vices block their own detection: it is immensely difficult for the closed-minded person to understand that she is closed-minded, and that is *because* of her closed-mindedness. Nevertheless, although it is difficult to detect stealthy vices it is not impossible, with traumatic experiences being one possible way of becoming aware of one's own stealthy vices.

With an account of both echo chambers and epistemic vices in place, we moved on to discussing how echo chambers can be seen through the lens of epistemic vices. We saw echo chamber members can be (and often are) guilty of several epistemic vices. Closed-mindedness, dogmatism, prejudice, and epistemic overconfidence are some of the epistemic vices that we find within some echo chambers. So, not so controversially, we found out that there are echo chambers which are epistemically vicious.

After understanding how epistemically vicious echo chambers look like, it then seemed appropriate to consider the possibility of echo chambers that are not epistemically vicious. After considering the example of the epistemic behavior of members of a university institute, we saw that there might be a case to be made for claiming that there is such a thing as a non-vicious echo chamber. If the justification for actively discrediting relevant voices in fact is sound, it might be wrong to claim that those who discredit are guilty of stealthy vices such as closed-mindedness and epistemic overconfidence. If what at first hand looks like closed-mindedness and epistemic overconfidence in fact does not serve as an obstacle to knowledge, then the epistemic behavior described is not vicious. Granted that our example of the university institute holds, we have a case for the existence of non-vicious echo chambers.

Considering that there is a distinction to be made between vicious and non-vicious echo chambers, we ended up suggesting a new definition of echo chambers which took this into consideration. This new definition placed importance at how the epistemic community in questions justifies its exclusion policy. Thus, we defined echo chambers as those communities which exclusion is justified in an epistemic vice. With this definition we are able to make a terminological distinction between vicious and non-vicious exclusion-based communities. Hopefully this specification can contribute something to future discussions on echo chambers.

Further research

Finally, before concluding this thesis, I would like to briefly discuss the potential avenues for further research on the topic of echo chambers.

First, as already mentioned, this thesis has had a focus on understanding epistemic vices as a collective concept, rather than examining the individual vices in greater detail. For example, while we have discussed the concept of prejudice, the assumed understanding of this particular vice has been somewhat rudimentary and has not taken into account the more nuanced and complex discussion of prejudice as described in works such as Endre Begby's *Prejudice* (2021). In order to more fully understand the nature and implications of epistemically vicious behavior, it would be beneficial to conduct a deeper investigation of the individual epistemic vices, and to explore the specific characteristics and challenges associated with each one. By taking a more detailed look at the individual epistemic vices, we can gain a better understanding of the ways in which they can impact our ability to acquire, evaluate, and use knowledge, and can develop more effective strategies for addressing and mitigating their effects. Given a vice-based understanding of echo chambers, a better understanding of the individual epistemic vices can also contribute to getting a better grasp of what echo chambers are, how they function and how they come to be.

Second, my inquiry has had a focus on echo chamber members, that is, I have looked at the epistemic characteristics of individual members of echo chambers. Although understanding the individual agents is vital to understand how echo chambers function, one cannot get a full grasp of echo chambers without analyzing how they function as a collective. Considering echo chambers through a more social epistemic lens might uncover important information, such as a deeper understanding of how echo chambers learn and exclude as a group, and how chamber-internal epistemic power structures function. Our epistemic agency is not something we necessarily conduct by ourselves, and social epistemic inquiry of echo chambers might uncover interesting facts and challenges posed by echo chambers being a collective endeavor.

Thirdly, the revised definition builds on the distinction made between those social epistemic structures which actively discredits relevant voices on the basis of epistemic vices and those structures that bases their discreditation in something else than epistemic vices. The former is

what I call an echo chamber, but what are we then to make of the latter? Our distinction invites a deeper inquiry of these non-vicious epistemic structures. Developing good and epistemically virtuous filtering techniques for processing information can be of great use for any person or institution that wishes to gain trustworthy knowledge, and a deeper inquiry within these non-vicious epistemic structures might unveil some important techniques.

Literature

Begby, E. (2021). *Prejudice*. Oxford University Press.

Blackwell Publishing Ltd. (2009). Dogmatism. In *The Blackwell Dictionary of Western Philosophy* (p. q90).

Boyd, K. (2019). Epistemically Pernicious Groups and the Groupstrapping Problem. In *Social Epistemology* 33(1), 61-73. doi:10.1080/02691728.2018.1551436

Cassam, Q. (2019). *Vices of the Mind*. Oxford University Press.

Thi Nguyen, C. (2020). Echo chambers and epistemic bubbles. In *Episteme*, 17(2), 141-161. doi:10.1017/epi.2018.32