A group identification account of collective epistemic vices

Kengo Miyazono[1] · Rie Iizuka[2]

Abstract

This paper offers an account of collective epistemic vices, which we call the "group identification account". The group identification account attributes collective epistemic vices to the groups that are constituted by "group identification", which is a primitive and non-doxastic self-understanding as a group member (Turner, 1982; Brewer, 1991; Brewer & Gardner, 1996; Pacherie, 2013; Salice & Miyazono, 2020). The distinctive feature of the group identification account is that it enables us to attribute epistemic vices not just to established social groups (e.g. committees, research teams, juries) but also to loose social groups (e.g. loosely connected people in an echo chamber) when they are constituted by group identification. The group identification account is contrasted with Fricker's (2010, 2020) influential account, the "joint commitment account", which focuses on established social groups, and has difficulty in making sense of collective epistemic vices of loose social groups. The group identification account is motivated by the fact that collective epistemic vices can be useful to diagnose not only the epistemic performance of established groups but also that of loose groups in real-life cases, such as echo chamber (Nguyen, 2020), implicit bias (Holroyd, 2020), group polarization (Broncano-Berrocal & Carter, 2021), etc.

Keywords Vice epistemology · Collective epistemic vices · Group identification · Virtue epistemology · Social epistemology

✉ Kengo Miyazono
miyazono@let.hokudai.ac.jp

Rie Iizuka
r201061@kansai-u.ac.jp

[1] Faculty of Humanities and Human Sciences, Hokkaido University, Nishi 7 Chome, Kita 10 Jo, Kita-ku, Sapporo Hokkaido 060-0810, Japan

[2] Faculty of Informatics, Kansai University, 2-1-1 Ryozenji-cho, Takatsuki Osaka 569-1095, Japan

1 Introduction

This paper offers an account of collective epistemic vices, which we call the "group identification account" ("GIA" hereafter). GIA attributes collective epistemic vices to the groups that are constituted by "group identification", which is a primitive and non-doxastic self-understanding as a group member (Turner, 1982; Brewer, 1991; Brewer & Gardner, 1996; Pacherie, 2013; Salice & Miyazono, 2020). The distinctive feature of GIA is that it enables us to attribute epistemic vices not just to established social groups (e.g. committees, research teams, juries) but also to loose social groups (e.g. loosely connected people in an echo chamber) when they are constituted by group identification.

Here is the background of our project. Collective virtue/vice epistemology[1] has both theoretical and applied aims. The theoretical aim is to answer theoretical questions about collective virtues/vices; about whether collective virtues/vices really exist, about what conditions need to be met for collective virtues/vices to be attributed, about whether collective virtues/vices are reducible to virtues/vices of the individuals, etc. The applied aim, in contrast, is to diagnose epistemic performance of groups in real-life cases[2] from the perspective of collective virtue/vice epistemology and explore possible prescriptions[3]. For instance, Fricker's influential work (Fricker, 2010, 2020) on collective virtue/vice epistemology includes both theoretical and applied parts. In the theoretical part, she presents a sophisticated theoretical account of collective virtue/vice; collective virtues/vices are grounded in "joint commitment" (Gilbert, 1989, 2013), which is an act of making clear to other members that one is willing to join forces with them in achieving some collective goal (see Sect. 2 for more about joint commitment). In the applied part, she uses her joint commitment account ("JCA" hereafter) of collective virtues/vices in order to diagnose epistemic performance of groups in real-life cases, such as institutional racism in London Metropolitan police (Fricker, 2010), or predatory sexual crimes at BBC (Fricker, 2020).

Fricker's JCA focuses on collective virtues/vices of established social groups (e.g. committees, research teams, juries, London Metropolitan police, BBC, etc.) that are typically constituted by joint commitment. Fricker's project is consistent with the observation that in everyday contexts collective virtues/vices are typically attributed to established social groups; e.g. "the research team is open-minded", "London Metropolitan police is racist", etc. But Fricker's focus on established groups can also be a limitation of her JCA, in particular in the applied part of collective virtue/vice epistemology. Collective virtues/vices can be useful to diagnose not only the epistemic performance of established groups but also that of loose groups in real-life

---

[1] For useful overviews, see Lahroodi (2018) and de Ridder (2022).

[2] de Rooij and de Bruin also talk about "the practical real-life relevance of an approach to collective virtue epistemology, which is also conceptually and empirically sound" (de Rooij & de Bruin, 2022, 397).

[3] Lahroodi also refers to the "applied collective virtue epistemology" which "investigates how epistemic virtues can be cultivated, maintained, or enhanced in groups and how the negative impact of individual members' epistemic vices on group epistemic functioning can be mitigated" (Lahroodi, 2018, 415). But our terminology of "applied" (where the emphasis is on *the application to real-life cases*) is different from and broader than Lahroodi's terminology (where the emphasis is on *the practical issues about cultivating, maintaining, or enhancing collective virtues*).

cases (see Sect. 2 for more about the distinction between established groups and loose groups). For instance, collective epistemic vices have been discussed in the context of explaining echo chamber (Nguyen, 2020), implicit bias (Holroyd, 2020), group polarization (Broncano-Berrocal & Carter, 2021), etc. where relevant groups are loose ones rather than established ones[4].

Fricker's JCA is not easily applicable to real-life cases where putative virtuous/vicious groups are loose groups because, unlike established groups, they are typically not constituted by joint commitment. For instance, as we will see in Sect. 2, JCA is not easily applicable to the cases of echo chamber. It is very unlikely that people in an echo chamber participate in a relevant kind of joint commitment; e.g. it is very unlikely that they make clear to others in the echo chamber that they are willing to join forces with them in achieving some collective goal (of, say, defending their idiosyncratic beliefs from counterevidence and objections).

There is thus a tension between the theoretical part and the applied part of collective virtue/vice epistemology. On the theoretical side, Fricker's JCA is philosophically sophisticated, and it is consistent with our everyday ascription of epistemic virtues/vices to established groups. On the applied side, in contrast, Fricker's JCA is not easily applicable to loose groups in real-life cases. A possible solution would be to bite-the-bullet; i.e. admit that JCA is not applicable to loose groups but insist that this is a plausible (or at least acceptable) consequence. For instance, one might admit that JCA is not applicable to a loose group in an echo chamber but insist that this is a plausible (or at least acceptable) consequence; e.g. because it is the wider social structure, not groups, that is epistemically problematic in the cases of echo chamber. The second option is to insist that JCA is in fact applicable to loose groups. For instance, one might argue that JCA, properly understood or suitably revised, is applicable to a loose group in an echo chamber. The third option is to propose an alternative account of collective epistemic virtues/vices that is clearly applicable to loose groups[5]. For instance, one might argue that there is an alternative account that is clearly applicable to a loose group in an echo chamber.

Our project in this paper is to explore the third option. Our proposal, GIA, which is an alternative to Fricker's JCA, is clearly applicable to loose groups such as the group in an echo chamber. Our discussion will proceed as follows. We give a brief overview of JCA and discuss the apparent difficulty of applying JCA to real-life cases of loose groups (Sect. 2). Then, we propose GIA and explain the main differences between JCA and GIA, and between "joint commitment" and "group identification" (Sect. 3). As a case study, we discuss echo chambers and show that GIA is applicable to loose groups in echo chambers and therefore that the loose groups can be described

---

[4] For this reason, Holroyd (2020) is dissatisfied with Fricker's JCA. Interestingly, however, Broncano-Berrocal and Carter (2021) appeal to a (reliabilist) version of JCA, which might be due to the fact that group polarization can happen between established groups such as scientific research teams. In fact, group polarization in science is their central example.

[5] For instance, Holroyd (2020) explores an alternative to JCA, motivated by real-life cases of implicit biases in loose groups. A similar project can be found in de Rooij and de Bruin (2022), although their focus is on large organizations rather than loose groups. They explore an alternative to JCA, motivated by real-life cases of large organizations such as Boeing.

as vicious (Sect. 4). In the last section (Sect. 5) we respond to the objection that GIA is committed to a problematic form of vice charging.

Some clarifications are in order.

First, we said above that our project in this paper is to explore the third option, i.e. presenting an alternative to JCA. It is not our claim, however, that the first option (i.e. simply admitting that JCA is not applicable to loose groups) and the second option (i.e. arguing that JCA is indeed applicable to loose groups) are hopeless. It is not part of our aim to show that the first and the second options do not work; rather our aim is only to show that the third option is at least a plausible option because there is indeed an alternative account that is clearly applicable to loose groups.

Second, GIA is presented as an "alternative" to JCA. Strictly speaking, these accounts might not be "alternatives" in a strong sense; i.e. they are incompatible with each other and hence at least one of them must be false. JCA can be interpreted as providing (jointly) sufficient conditions for collective virtues/vices rather than necessary conditions[6]. Understood in this way, JCA (which does not say that joint commitment is necessary) is compatible with GIA (which says that group identification is sufficient, and is neutral, at least in principle, on the necessity of joint commitment)[7].

2 Joint commitment account

2.1 Joint commitment account

Fricker's JCA (Fricker, 2010, 2020; see also Kidd, 2021) is based on Gilbert's (1989, 2013) theory of plural subjects. The core idea of JCA is that collective virtues/vices are grounded in what Gilbert calls "joint commitments". JCA says that, roughly, a group of individuals, G, is collectively virtuous/vicious when the individuals, operating under the practical identity as the members of G, participate in joint commitment of a right kind.

Joint commitment: An individual, X, participates in joint commitment with the other members of G (e.g. a poetry reading group) by making clear to them, under the condition of common knowledge, that X is willing to join forces with them in accepting a collective goal (e.g. finding the best interpretation of a poem). An important feature of joint commitment is that an explicit statement of agreement is not necessary in order for X to participate in joint commitment and make clear to others that X is willing to join forces with them. X can participate in joint commitment in an implicit

---

[6] For instance, Fricker (2010, 2020) does not rule out summative cases of collective virtues/vices where collective virtues/vices are attributed to G because all (or at least most) members of G have the virtues/vices. The distinction between summativism and non-summativism "should not be thought of as a competition, for both models represent perfectly real possibilities, and indeed each kind of group is frequently realized in institutional life" (Fricker, 2020, 93). This suggests that the conditions in JCA are not intended as necessary for collective virtues/vices.

[7] This allows for the possibility that there are some cases of collective virtues/vices that are explained by JCA but not by GIA; e.g. the case of a vicious established group where members, without group identification, reluctantly participate in the relevant joint commitments under coercion. We thank an anonymous referee for raising this issue.

way, by letting it stand, or going along with an idea (e.g. a reading group member participates in joint commitment by going along with an interpretation during the discussion). Another important feature of joint commitment is that, once X participates in a joint commitment, X is subject to rebuke or at least demand for explanation when X fails to fulfill the joint commitment (e.g. a reading group member openly rejects the official conclusion of the group[8]).

The "right kind of" joint commitment is different in different cases of collective virtues/vices. In motive-based cases, the right kind of joint commitment is the one to virtuous motives[9]. Fricker's example is a research team whose members jointly commit to the motives of diligence and thoroughness. In skill-based cases, the right kind of joint commitment is the one to virtuous ends. Fricker's example is a night watch team whose members jointly commit to the end of staying on the alert for enemy movement by way of an efficient diffusion of labor.

Practical Identity: When X participates in joint commitment, X does so as a member of G rather than as a private individual. Fricker expresses this in terms of the concept of "practical identity", which she inherits from Korsgaard (1996): X participates in joint commitment under the practical identity as a member of G. X's attributes under the practical identity as a member of G can come apart from X's attributes as a member of another group F, and from X's attributes as a private individual. For instance, it is possible that a person cares about the value of diverse opinions as a member of the poetry reading group, but she does not care about it as the CEO of an IT startup, or as a private individual.

This is related to one of the distinctive features of JCA; i.e. it is a non-summative account of collective virtues/vices. Roughly, summativism is the view that collective virtues/vices are attributed to G only if all (or at least most) members of G have the virtues/vices. Non-summativism denies this. JCA is a non-summative account because, as we just noted, X's attributes under the practical identity as a member of G can come apart from X's attributes as a private individual. In particular, it is possible that people in G participate in joint commitment of a right kind under the practical identity as the members of G (such that G is regarded as collectively virtuous/vicious), without having the virtues/vices as private individuals. For instance, people in a poetry reading group participate in joint commitment of a right kind under the practical identity as the members of the reading group (such that the reading group is regarded as collectively open-minded), without being open-minded as private individuals.

2.2 Joint commitment and collective epistemic vices

As we noted, collective virtue/vice epistemology has both theoretical and applied aims. The theoretical aim is to answer theoretical questions about collective virtues/

---

[8] However, this does not prevent the reading group member from having a different opinion *personally* rather than *publicly*. She can even express a different opinion with preamble, making "it clear that one is speaking for oneself alone, and not for the group" (Gilbert, 1987, 199). For instance, she can say "Still I personally think that the interpretation is incorrect" without being subject to rebuke.

[9] Strictly speaking, the joint commitment to a virtuous motive is not sufficient for motive-based virtues. In addition, a reliability condition of some kind needs to be satisfied.

vices. The applied aim, in contrast, is to diagnose epistemic performance of groups in real-life cases from the perspective of collective virtue/vice epistemology and explore possible prescriptions.

When it comes to the theoretical aim, the most sophisticated theoretical account of collective virtues/vices is Fricker's JCA, whose core idea is that collective virtues/vices are grounded in joint commitment. JCA focuses on established social groups (e.g. committees, research teams, juries, etc.), which are typically constituted by joint commitment. In fact, Fricker's project is motivated by the statements in everyday discourse where epistemic virtues/vices are attributed to established social groups: e.g. "The jury is fair-minded."

When it comes to the applied aim, however, JCA's focus on established groups can be a limitation. Collective virtues/vices can be useful to diagnose not only the epistemic performance of established groups but also that of loose groups in real-life cases[10]. According to Lahroodi, what distinguishes established groups from loose groups (or, what Lahroodi calls "loose associations") is that "the former are relatively coherent units in which members are bounded and united together in some fashion, there is a high degree of interaction among members, and consequently such groups are capable of action in a manner not dissimilar to that of a single subject or agent" (Lahroodi, 2018, 407–408).

Think, for example, about Nguyen's (2020) suggestion that echo chambers are collectively vicious (see Sect. 4 for details). Typically people in an echo chamber do not form an established group; they are a group in a loose sense. It is very unlikely that loosely connected people in an echo chamber participate in joint commitment of a relevant kind. It is very unlikely that they will make it clear to others in the echo chamber that they are willing to join forces with them in achieving some collective goal (of, say, defending their idiosyncratic beliefs from counterevidence and objections)[11].

Before moving on, we briefly discuss a general ("general" in the sense that it is not peculiar to the context of applying JCA to loose groups) issue of JCA and its application to vices.

There is a disagreement among theorists about the role of motivation in epistemic vices. On the one hand, epistemic virtues require reliable dispositions of reaching epistemic goods out of a person's motivation toward such goods (Zagzebski 1996). A mere success is not sufficient for a virtue; a stable success from a good motive is necessary. On the other hand, the role of motivation in vices is unclear. Vices cannot simply be understood in terms of the presence of bad motivations or the absence of good ones (Crerar, 2018). We call people vicious not necessarily because of the

---

[10] As Medina points out, epistemic virtues/vices can be applied to "very diverse forms of groupings and social formations ranging from highly structured and rigidly organized groups (such as the Republican party) to diffused and changing networks or publics (such as Fox News' viewership)" (Medina, 2021, 337).

[11] Similarly, Holroyd's (2020) suggests that implicit biases constitute a vice at the collective level. Typically implicitly biased people in a community do not form an established group; they are a group in a loose sense: "It is entirely implausible to suppose that there is a joint commitment to some bad motive, or bad epistemic end, involved in cases where groups stably manifest implicit biases – that each participant of the loosely connected group has committed to make discriminatory judgments about the value of women philosophers, say, or to ignore contributions, or dismiss lines of argument" (Holroyd, 2020, 136).

presence of bad motivations or the absence of good motivations but rather because these people have no epistemic ends (e.g. extreme indolence), because they are genuinely motivated by an ultimate epistemic end yet exemplify vices (e.g. Galileo's arrogance), etc. Hence, vices lack "a psychological unity" (Crerar, 2018, 762); there are multiple ways to fall prey to vices.

Fricker's JCA retains the motivationalist spirit (Fricker, 2010, 2020)[12]. She characterizes vices as culpable epistemic failures of epistemic virtue in motivations (or *ethos*), implementation, or both (Fricker, 2020). Fricker admits that some may be vicious while genuinely motivated by ultimate epistemic goods because of their failures at the level of their performances or because of their persistent failures regarding the motivation for the mediate end (e.g. a failure of fact-checking while being motivated to acquire knowledge). Still, a question remains as to how vices can be distinguished from non-virtuous neutral traits in Fricker's theory. Some failures of (ultimate/mediate) motivation or performance are plausibly regarded as vices, but it is not the case that all the failures of virtues amount to vices. A man striving for knowledge and hence trying to be intellectually humble may occasionally fail to act according to his motivation. It depends on the details of the case whether such a trait is a vice or a non-virtuous trait. Fricker's theory, as it stands, may overmultiply vices and have difficulty in distinguishing vices from non-virtuous neutral traits (Holroyd, 2020).

In any case, this issue of JCA is a general one, which is about applying JCA to epistemic vices in general. Our focus is rather on an issue that is peculiar in the context of applying JCA to vicious loose groups; loosely connected people do not form the right kind of group in the first place.

3 Group identification account

Our project in this paper is to explore an account of collective epistemic vices that is clearly applicable to loose groups such as the one in an echo chamber. This section presents what we take to be a promising candidate; GIA[13].

Before presenting our proposal, we briefly discuss an obvious candidate, namely, the simple summative account. Perhaps attribution of virtues/vices to loose groups can be understood in a summative manner (in which case the problem is solved by Fricker's pluralism according to which summativism is true in some cases and non-summative JCA is true in other cases; see the footnote 6).

The simple summative account, however, is not appropriate for our purpose because it does not seem to be applicable to the kind of cases we are interested in. For instance, it is far from obvious that the collective vice of echo chamber can be understood in a summative manner. It is far from obvious that individuals in an echo

---

[12] Broncano-Berrocal and Carter (2021) also raise the issue of the motivationalist commitment in Fricker's JCA.

[13] GIA is not the only candidate. Other accounts that are based on similar motivation include Holroyd's (2020) account, which is based on Byerly & Byerly's (2016) account, and de Rooij & de Bruin's (2022) account.

chamber are epistemically vicious (Nguyen, 2020; see also Medina, 2021; Begby, 2022; Levy, 2022b). Rather, according to Nguyen, an echo chamber turns individually virtuous people into a collectively vicious group[14]. See Sect. 5 for more details.

Thus, what we need is a non-summative account that is clearly applicable to loose groups such as the group of loosely connected people in an echo chamber. Now we are ready to present our own proposal, GIA, which is summarized as follows: a group of individuals, G, is collectively vicious when (1) the individuals group-identify with G and (2) thereby acquire motivation to act as the members of G such that (3) G is disposed to act in a vicious manner.

We will now explain the notion of group identification, which is central to GIA, as well as the similarities and differences between group identification and joint commitment.

Group identification: An individual "group-identifies" with G (or acquires the "group identification" as a member of G) when she achieves a primitive and non-doxastic form of self-conception as a member of G[15].

It has been suggested in social psychology and philosophy that group identification plays a crucial role in collective, cooperative, and collaborative action and cognition (Turner, 1982; Brewer, 1991; Brewer & Gardner, 1996; Pacherie, 2013; Salice & Miyazono, 2020; Miyazono & Inarimori, 2021; Pickard, 2021).

Experimental studies, in particular the ones in the minimal group paradigm (for an overview, see Diehl, 1990), show that group identification motivates pro-group behavior. In a classic study by Tajfel et al. (1971), for example, a person with the group identity as a Kandinsky-lover (as opposed to a Klee-lover) was motivated to be helpful to other Kandinsky-lovers (as opposed to Klee-lovers). In the experiment, participants first indicated their aesthetic preferences in response to given pairs of paintings. Based on their preferences, they were told that they were either Kandinsky-lovers or Klee-lovers (while, unbeknownst to them, they were randomly assigned to either the Kandinsky-lover group or the Klee-lover group). In the next task, participants were asked to distribute real monetary rewards to other participants. During this task, in-group favoritism was observed: Kandinsky-lovers favored other Kandinsky-lovers and Klee-lovers favored other Klee-lovers.

---

[14] Similarly, Holroyd (2020) argues that implicit biases cannot be understood as epistemic vices of individuals. The scores on implicit measures are unstable, and they are not systematically correlated with personal level discriminatory behavior. Thus, at the individual level, implicit biases do not meet the condition that epistemic vices systematically obstruct inquiry or the condition that epistemic vices reliability produce false beliefs. Rather, implicit biases have systematic and stable discriminatory effects at the situational and contextual level (Payne et al., 2017), which suggests that implicit biases can be understood as an epistemic vice at the collective level in a non-summative sense.

[15] Battaly's (2022) account of solidarity as a distinctively collective virtue has the condition that individuals feel the sense of belonging to the group, which is similar to group identification. In fact, her account of solidarity is similar to GIA in several ways such that the former can be cashed out by the latter. According to Battaly, a group G has the trait of solidarity to the extent that its members are disposed to (1) share values, aims, or goals; (2) care about the values, aims, or goals; (3) act in accordance with the values, aims, or goals; (4) trust the testimony of other group members with respect to the values, aims, and goals; and (5) feel a sense of belonging to the group.

Group identification plays a central role in Pacherie's (2013) "lite" theory of shared intention and joint action. According to the lite theory, two persons P1 and P2 share an intention to A, if:

(i)   each has a self-conception as a member of the team T, consisting of P1 and P2;
(i') each believes (i);
(ii) each reasons that A is the best choice of action for the team;
(iii) each therefore intends to do his part of A.

(i) amounts to the condition that each acquires group identification as a member of T. The lite theory is "lite" in comparison to competing theories of joint action and shared intention, such as Bratman's (2009) theory, for several reasons. First, the lite theory allows for joint action and shared intention in situations where agents cannot communicate with each other. For instance, agents in a one-shot version of the Prisoner's Dilemma cannot communicate with each other, but they might be able to acquire group identification and thus satisfy the condition (i) (Bacharach, 2006). Second, the lite theory makes joint action and shared intention less cognitively demanding than competing theories. This is due to the fact that group identification is a primitive form of self-understanding, which does not require sophisticated social cognitive capacities. For instance, infants without sophisticated social cognitive capacities can acquire group identification and thus satisfy the condition (i).

How does group identification work, exactly? One possibility is that group identification is a form of framing; people identify themselves as the members of a group when they adopt a group frame to represent themselves and the agents with whom they are interacting (Bacharach, 2006). In this case, group identification's influence on act and cognition is a special case of framing effect (Bermudez, 2022).

Salice and Miyazono (2020) propose a representational account of group identification, according to which group identification is cashed out in terms of non-doxastic representations. They reject the idea that group identification is a doxastic process in which, for example, X comes to believe that X is a member of G; such a belief does not account for the motivational force of group identification. Rather, group identification involves a "pushmi-pullyu representation" (Millikan, 1995, 2004): when X group-identifies with G, X forms a representation with both descriptive content (e.g. describing oneself as a member of G) as well as directive content (e.g. directing oneself to behave as a member of G).

Pushmi-pullyu representations are representational states that are evolutionarily and structurally primitive, that have both descriptive content and directive content, and that are intrinsically motivating without being combined with a conative state. One of Millikan's examples is the waggle dance of honey bees. Honey bees in the same colony do the "dance" to share the information about important locations such as the location of flowers yielding nectar and pollen. The dance is a pushmi-pullyu representation; it does the job of describing the location of flowers as well as the job of directing the behavior of flying to the location.

The waggle dance is triggered in the presence of relevant cues (e.g. flowers), and it initiates the relevant behavioral processes (e.g. another bee's flying to the flowers). Group-identification is also triggered by some cues. Social psychologists have

already identified several "group cues" (Pacherie, 2013) that are able to trigger group identification. For instance, in his survey of the literature, Bacharach (2006) discusses belonging to the same social category, having common interests, sharing common fate, facing a competing group, employing we-language etc. as properties in the environment that, once they acquire salience for the subject, may cause group identification.

Group Identification and Joint Commitment: Let us now compare group identification, which is central to GIA, and joint commitment, which is central to JCA.

Group identification can motivate or facilitate cooperative activity without joint commitment. An individual X can group-identify with G and thereby become motivated to act as a member of G without joint commitment with other members of G or, in other words, without making clear to other members of G that X is willing to join forces with them in accepting collective goals. In the Kandinsky-Klee experiment, for example, a person with the group identity as a Kandinsky-lover, let us call her Ruth, was motivated to behave as a Kandinsky-lover without joint commitment with other Kandinsky-lovers or, in other words, without making clear to other Kandinsky-lovers that she is willing to join forces with them in accepting collective goals.

This is related to another difference between group identification and joint commitment, which is crucial for our project in this paper. Joint commitment typically requires an established group; typically an individual X participates in a joint commitment with other members of an established group G. Group identification, in contrast, does not require established groups; an individual X can group-identify with G and thereby become motivated to act as a member of G even if G is not an established group. In the Kandinsky-Klee experiment, for example, Ruth was motivated to act as a Kandinsky-lover where Kandinsky-lovers are not an established group. Because of this feature of group identification, GIA is applicable not only to established groups but also to loose groups.

Here is another possible difference between group identification and joint commitment. As we already indicated above, once X participates in a joint commitment (e.g. to value diverse opinions) under the practical identity as a member of G (e.g. the poetry reading group), X is subject to rebuke when X fails to fulfill the joint commitment (e.g. fails to value diverse opinions in a reading group meeting). In contrast, it does not seem to be the case that, once X group-identifies with G (e.g. Kandinsky-lovers) and is motivated to act as a member of G (e.g. help other Kandinsky-lovers), X is subject to rebuke when X suddenly changes her mind and group-identifies with another group (e.g. Klee-lovers). This suggests that, in a sense, group identification is less stable than joint commitment. One might think that this is a problem for GIA; e.g. the *instability* of group identification makes it difficult for GIA to account for (collective) virtues/vices that are typically conceived as *stable* dispositions[16]. However, the instability of this kind only shows that an individual, X, can change her group-identification, which is compatible with the idea that the group, G, has a stable disposition. For instance, even when Ruth does not identify herself as a Kandinsky-lover any longer, Kandinsky-lovers (with which Ruth does not group-identify) can still have the stable disposition to be helpful toward other Kandinsky-lovers.

---

[16] We thank an anonymous referee for raising this issue.

In addition to the differences above, there is a similarity between group identification and joint commitment; both are non-summative. The non-summative nature of JCA stems from the fact that when an individual X participates in a joint commitment, X does so, not necessarily as a private individual, but rather under *the practical identity as a member of G*. For instance, a person might jointly commit to the value of diverse opinions as a member of a poetry reading group, not as a private individual. Perhaps she does not care about it in other contexts; e.g. when she is behaving as the CEO of an IT startup. Similarly, the non-summative nature of GIA stems from the fact that when an individual X is motivated by group identification, X is so motivated, not necessarily as a private individual, but rather under *the group identity as a member of G*. For instance, Ruth was motivated to help other Kandinsky-lovers in the experiment as a Kandinsky-lover, not as a private individual. Perhaps she is not very helpful to Kandinsky-lovers in other contexts in which another group-identification (e.g. group identification as a Yale alumnus) is salient.

One might think, however, that GIA is summative *in a sense*; collective vices according to GIA are to be explained by the sum of a particular psychological state (i.e. the state of group identification) of individuals. In Fricker's phrases, GIA is compatible with "the individualistic spirit of summativism" where collective vices are "to be understood as none other than the sum of the individual's contributions" (Fricker, 2020, 237). We do not deny this, but note that GIA is still non-summative *in the sense we indicated above*; i.e. it is not the case, according to GIA, that a vice is attributed to G only if all (or at least most) members of G have the vice. When a vice is attributed to G, it is assumed, according to GIA, that each individual member group-identifies with G, and is motivated to behave as a member of G, but this does not imply that the members of G have the vice individually; i.e. the blameworthy character trait that systematically obstructs inquiries might not be attributed to the individual members of G. As we will see in the next section, GIA allows for the possibility that a vice is attributed to a group of loosely connected people in an echo chamber while the vice is not attributed to the individual members.

4 Case study: echo chambers

Unlike JCA, GIA can attribute epistemic vices to a loose group G when G is constituted by group identification; i.e. when the members of G acquire the self-conception as group members and become motivated to act as such. This section discusses the phenomenon of echo chamber as a case study and shows that GIA can attribute collective epistemic vices to loosely connected people in an echo chamber.

We follow Nguyen's influential distinction between epistemic bubbles and echo chambers. An epistemic bubble is "a social epistemic structure in which some relevant voices have been excluded through omission", and "can form with no ill intent, through ordinary processes of social selection and community formation" (Nguyen, 2020, 142). In contrast, an echo chamber is "a social epistemic structure in which other relevant voices have been actively discredited", and "work[s] by systematically isolating their members from all outside epistemic sources" (Nguyen, 2020, 142).

The distinction between epistemic bubbles and echo chambers matters when it comes to the attribution of epistemic vices. It is relatively easy to attribute epistemic vices to individuals in epistemic bubbles. As Nguyen points out, people in epistemic bubbles fail to incorporate some information often due to their epistemic laziness; "if you're subject to an epistemic bubble because you get all your news from Facebook and don't bother to look at other sources, you are, indeed, blameworthy for that failure. If one finds the language of epistemic virtues and vices appealing, then we can say that members of epistemic bubbles are committing the vice of epistemic laziness" (Nguyen, 2020, 154).

In contrast, it is difficult to attribute epistemic vices to individuals in echo chambers. Nguyen discusses a case of a teenager who has grown up among family members and close friends who are firmly committed to an echo chamber. Trapped in the echo chamber, the teenager acquires false or irrational beliefs, but perhaps there is nothing epistemically vicious about him: "Our teenager seems, in fact, to be behaving with many epistemic virtues. They are not at all lazy; they are proactive in seeking out new sources. They are not blindly trusting; they investigate claims of epistemic authority and decide for themselves, using all the evidence and beliefs that they presently accept, whether to accept or deny the purported expertise of others. They have theories, which they have acquired by reasonable methods, predicting the maliciousness of outsiders; they increase their trust in those theories when their predictions are confirmed" (Nguyen, 2020, 155)[17].

Thus, if echo chambers involve epistemic vices, the vices are to be found in the collective level rather than in the individual level, and the epistemic vices in the collective level cannot be summative. Nguyen suggests that "echo chambers are social epistemic structures which convert individually epistemically virtuous activity into collective epistemic vice (Nguyen, 2020, 155)[18]. An echo chamber is analogous to, but also contrasted with, Mandevillian intelligence (Smart, 2018) where individually epistemically vicious activity is converted into collective epistemic virtue; e.g. academia constitutes a structure where intellectual stubbornness at the individual level is converted into the intellectual carefulness and thoroughness at the collective level[19].

---

[17] Similarly, Levy (2022b; see also Levy 2022a) argues that echo chamber is a modern version of a familiar and reasonable pattern of belief formation in which we selectively place more weight to testimony of in-groups rather than out-groups: "Preferring the testimony of the in-group to that of the out-group therefore seems justifiable. In trusting my fellow partisans, I trust those who tend to get things right in the factual and the normative domains. Conversely, in distrusting out-group testimony, I discard the testimony of those who are less likely to be competent and more likely to be deceptive. None of this is epistemically objectionable. […] Echo chambers cannot be objectionable on the grounds that they lead us to discount evidence that is preempted by our priors; such preemption is how rational agents are supposed to function" (Levy, 2022b, 7–8).

[18] A similar idea was suggested by Medina: "The epistemic vices of groups that become echo chambers are often not reducible to the individual attitudes and habits of their members, but have to be understood as epistemic group traits that are sustained by the epistemic and discursive norms of the group, that is, by the epistemic-communicative ecosystem in which the group operates" (Medina, 2021, 339).

[19] One might think, however, that Nguyen's account of echo chambers is also compatible with the idea that it is the hostile epistemic environment or structure that is vicious or problematic; people are not (collectively) vicious at all. It is not obvious, however, that this is the best interpretation of Nguyen's account. First, social epistemic structures are certainly crucial in Nguyen's account, but their role is to "convert individually epistemically virtuous activity into collective epistemic vice" (Nguyen, 2020, 155). Second,

JCA has difficulty in making sense of the idea of echo chamber as being vicious at the collective level. Typically people in an echo chamber are only loosely connected in such a way that it is extremely unlikely that they participate in joint commitment in Gilbert-Fricker's sense. It is extremely unlikely, for example, that the teenager in Nguyen's example makes clear to others in the echo chamber that he is willing to join forces with them in achieving some collective goal.

In contrast, GIA has a better chance for making sense of the idea of echo chamber as being vicious at the collective level. The central hypothesis here is that people in an echo chamber group-identify, which motivates them to act and think as group members[20]. People in an echo chamber group-identify with a particular group that is associated with the echo chamber[21] and thereby become motivated to act and think as group members; e.g. become motivated as group members to defend their own view (which can be associated with their group identity) and discredit opposing views[22]. This idea is consistent with Nguyen's characterization of echo chamber as "an epistemic community which creates a significant disparity in trust between members and non-members" where "this disparity is created by excluding non-members through epistemic discrediting, while simultaneously amplifying members' epistemic credentials" (Nguyen, 2020, 146).

Nguyen's distinction between epistemic bubbles and echo chambers matters here again. On the one hand, epistemic bubbles have little (or nothing) to do with motivation as a group member. Think, for example, about so-called "filter bubbles" that are created by algorithmic filtering. Beliefs of those in a filter bubble are biased because of the omission of some information due to algorithmic filtering, not because of the motivation to defend their own view and discredit opposing views. On the other hand, echo chambers work differently, not by the mere omission of some information but rather by active discrediting some (source of) information. Echo chambers are "more malicious" than epistemic bubbles, and they "are set up intentionally, or at least maintained for this functionality" (Nguyen, 2020, 149). It is plausible, then, to think that, unlike epistemic bubbles, echo chambers have something to do with moti-

---

as we already noted, Nguyen compares echo chambers with Mandevillian intelligence (where collective epistemic virtues arise out of epistemically vicious activities of individuals). This contrast suggests that, in echo chambers, collective epistemic vices arise out of epistemically virtuous activities of individuals.

[20] Thus, this account of echo chamber is broadly a motivational one (cf. Avnur, 2020).

[21] Anderson's discursive model of epistemic bubbles also focuses on the role of group identity; "Identity-expressive discourse expresses the speaker's group identity, and positions the speaker in relation to people with the same or other identities. […] a group adopts discursive norms to treat certain ostensibly empirical assertions as identity-expressive. It thereby removes those utterances for empirical inquiry, in which assertibility is governed by evidence. A group following such norms may thus create the functional equivalent of an epistemic bubble, by blocking accuracy-guided empirical inquiry within the group with respect to these assertions" (Anderson, 2021, 23).

[22] Motivation of this kind can easily be found in everyday life and also in academic discourse. An example from academic philosophy can be found in Quine's homage to Carnap, in which Quine recalls his young day when he identified himself as a Vienna-style empiricist, was motivated to defend empiricist philosophy from objections by American philosophers, and was delighted when he heard about Carnap's response to Lovejoy's criticism of empiricist philosophy: "We beamed with partisan pride when [Carnap] countered a diatribe of Arthur Lovejoy's in his characteristically reasonable way, explaining that if Lovejoy means A then p, and if he means B then q. I had yet to learn how unsatisfying this way of Carnap's could sometimes be" (Quine & Carnap, 1991, 465).

vation to defend their own view and discredit other views. Beliefs of those in an echo chamber are biased not because of the mere omission of some information but rather because of group identification, which motivates group members to defend their own view and discredit opposing views.

As we already noted, group identification is triggered by "group cues" (Pacherie, 2013; Bacharach, 2006) including belonging to the same social category, having common interests, sharing common fate, facing a competing group, employing we-language, etc. These group cues can be found in echo chambers. For instance, common vocabulary and common enemies play a crucial role in echo chambers, according to Jamieson and Cappella (2008). Indeed, Jamieson and Cappella associate (conservative) echo chambers with group identity or collective identity that is built and sustained by a use of common vocabulary and by identifying common enemies: "Communities create a collective identity and the bonds that sustain it. They do so, in part, by employing a distinctive common vocabulary that carries with it a way of seeing the community and its adversaries. […] Creating a common enemy is […] a central means of establishing and sustaining a group identity. […] The disdainful language elicits emotion that creates bonds within the community and reinforces the notion that Democrats and their polity positions are the enemy" (Jamieson & Cappella, 2008, 179–180).

To sum up, GIA opens up the possibility of attributing epistemic vices to loose groups in echo chambers. Of course, there are remaining empirical issues as to how echo chambers work (e.g. whether they really involve group identification) and how people in them think and behave (e.g. whether they are really motivated as group members to defend their own view and discredit other views), which are relevant to whether collective epistemic vices are really attributed to loose groups in echo chambers (according to GIA). We need more empirical data to address these issues.

5 Vice charging objection

5.1 Vice charging

Before closing, we will address an issue about GIA and its application to loose groups. So far, we have been exploring an account of collective vices that enable us to attribute collective epistemic vices to groups of loosely connected people. We argued that, on the one hand, JCA does not enable us to attribute epistemic vices to loose groups and, on the other hand, GIA has a better chance, which we take to be an advantage of (or, at least, an attractive feature of) GIA.

A possible objection to our argument is that attributing epistemic vices to loose groups is not an advantage of GIA because we should not attribute epistemic vices to loose groups in the first place. Indeed, our project of attributing epistemic vices not only to established groups but also to loose groups might be regarded as a problem-

atic form of "vice charging" (Kidd, 2016)[23] where attribution of epistemic vices is excessive (Cassam, 2021a, b).

Before responding to this vice charging objection, we need to be precise about what the objection actually is. We can distinguish two possible vice charging worries; the worry that the very idea of attributing epistemic vices to groups (rather than individuals) is committed to an excessive attribution of epistemic vices; and the worry that attributing epistemic vices to loose groups (rather than established groups) is committed to an excessive attribution of epistemic vices.

Fricker briefly discusses the former worry (or something similar) in the context of the discussion of institutional vice: "Why might we care about the question whether institutions can be said to have vices of any kind? Why not content ourselves with using a vocabulary of, say, efficiency and inefficiency, functionality and dysfunctionality, relative to the institution's goals or purposes?" (Fricker, 2020, 89). Fricker's response is that the idea of institutional vice (or something similar) can be found in the public discourse (e.g. the London Metropolitan Police was described as "institutionally racist"), and the idea can serve the distinctive purpose "of picking out aspects of institutions that are the collective analogue of an individual agent's character, but where the actual individuals whose combined epistemic agency comprises the institution's epistemic agency need not, as individuals, have any of the traits or attitudes of the institution" (Fricker, 2020, 90).

Whether Fricker's response to the first worry is fully satisfactory or not, we do not discuss it here. The first worry is about collective vice epistemology in general, not about our project in particular. Our focus is rather on the latter worry, which can be rephrased as follows; even granting that the very idea of attributing epistemic vices to groups is not necessarily problematic, still attributing epistemic vices to loose groups is a problematic form of vice charging.

What does a problematic vice charging look like, exactly? Cassam (2021a) argues, for instance, that vice explanation of MMR vaccine hesitancy (where the vaccine hesitancy is explained by the epistemic vices of individuals) is committed to a problematic and vicious form of vice charging, failing to take into account relevant ethnographic and sociological factors, such as the culture of particularistic view of child health where each child is regarded as unique and deserving individualized health care. More generally, Cassam argues that attributions of epistemic vices to individuals can be problematic when one of the following conditions is satisfied:

1: Their assumption that the attributee's conduct is epistemically defective is open to question, especially in cases where this assumption is grounded in partisan political differences between the attributer and attributee.

---

[23] For Kidd (2016), however, "vice charging" refers to the actual practice in our society where people charge others with epistemic vices. Our problem here is not about the actual practice in our society, but rather about the theoretical project where epistemic vices are attributed (to groups of loosely connected people). One might say, then, that strictly speaking our issue is not a problem of vice charging per se but rather a problem of developing a theory of epistemic vices that *justifies* a potential form of vice charging where loosely connected people are charged with epistemic vices.

2: Even if the attributee's conduct is epistemically defective, there may be better ways of explaining its defectiveness than by pinning it on an underlying epistemic vice.

3: Vice attributions potentially underestimate the extent to which epistemically vicious thinking can nevertheless be rational. Even in epistemically vicious thinking there must be some semblance of cogency. (Cassam, 2021b, 302)

Our project of attributing epistemic vices to loose groups does not satisfy any of the conditions above. First, *individualist* vice epistemology might satisfy the condition 2, but *collective* vice epistemology does not. For instance, an individualist vice explanation of a phenomenon (e.g. MMR vaccine hesitancy) can be problematic when it ignores other relevant factors, such as social or situational factors. In contrast, collective vice explanation of a phenomenon does not ignore social or situational factors. In fact, it is one of the distinctive features of collective vice explanation that social or situational factors are incorporated in it; e.g., social and political factors are crucial ingredients of a plausible collective vice explanation of echo chamber. As Nguyen (2020) notes, echo chambers are *social epistemic structures* in which individually epistemically virtuous activities are converted into a collective epistemic vice.

Second, (some versions of) *summative* collective vice epistemology might satisfy the condition 1 or 3, but *non-summative* collective vice epistemology does not. For instance, a summative collective vice explanation of echo chamber can be problematic because it (wrongly) assumes that individuals in an echo chamber are epistemically vicious or irrational. In contrast, a non-summative collective vice explanation of echo chamber does not assume that individuals in an echo chamber are epistemically vicious or irrational. In fact, it is the distinctive features of non-summativism that we can attribute epistemic vices to a group without attributing them to its individual members.

Thus, at least non-summative collective vice explanation does not involve a problematic vice attribution to individuals. But one might still think that even non-summative vice attribution to groups can be problematic. Cassam's worries can arise even in the context of non-summative vice attribution to groups; e.g. the worry that it is not obvious that the relevant groups are epistemically defective, the worry that the epistemic defectiveness can be explained without attributing vices to the groups, etc. For instance, one might think that echo chambers are *entirely* due to social or situational misfortunes, and thus there is nothing epistemically defective or vicious about the groups.

At least according to Nguyen's characterization, however, it does not seem to be the case that echo chambers are *entirely* due to social or situational factors. Perhaps some cases of filter bubbles are *entirely* due to social or situational factors such as misleading algorithmic filtering (although, as Nguyen points out, one can still argue that people are epistemically lazy in failing to seek information outside filter bubbles). In contrast, echo chambers are not *entirely* due to social or situational misfortunes. For example, as we already noted in Sect. 4, echo chambers are characterized as "malicious" entities in where some information is "actively discredited" by their members who are motivated to defend their own view and discredit opposing views (see also the footnote 19).

But it is still possible that Nguyen's characterization is wrong and echo chambers in the real world are in fact entirely due to social or situational factors. For instance, an apparent malicious active discrediting of information may not actually be malicious nor active; some social or situational factors may explain them away. We do not have any empirical evidence that rules out such a possibility. This brings us to the limitation of this paper that we mentioned in the end of Sect. 4 (and will mention in the concluding section as well); there are open empirical issues concerning how echo chambers work and how people in them think and behave.

5.2 Too many vices?

Even if our project can be defended from the vice charging objection, however, it is still possible that some particular accounts of collective epistemic vices invite a serious worry of excessive attributing epistemic vices. For instance, Holroyd (2020) argues that collective epistemic vices of loose groups can be understood as invisible-hand vices (Fricker, 2010) in which the group-level disposition (e.g. of being implicitly biased collectively) emerges, by the invisible-hand mechanism, out of individuals who do not have corresponding individual-level disposition (e.g. of being implicitly biased individually). The invisible-hand account of collective epistemic vices faces the problem of excessive attributions; it is a very liberal account that attributes vices to groups, no matter whether they are united by joint commitment or group identification, even when there is no interesting sense in which people in the groups constitute a collective agency. In contrast, GIA can easily be linked to collective agency, e.g. via Pacherie's (2013) lite theory of shared intention and joint action.

This brings us to a difficult issue of how to distinguish a warranted attribution of epistemic vices and an excessive attribution of epistemic vices. We do not have a clear criterion, unfortunately. We will end this paper with some suggestions, referring to general accounts of vices by Cassam (2019) and Battaly (2016).

According to Cassam, epistemic vices are "blameworthy or otherwise reprehensible attitudes, character traits or ways of thinking that systematically obstruct the gaining, keeping or sharing of knowledge" (Cassam 2019, 1). Cassam does not require acquisition responsibility for vices. What matters for vices is that the relevant trait is malleable by the hands of the agents for the future, not that the trait is acquired due to the agents' will or choices in the past. If characters, attitudes, or ways of thinking can be "actively reconfigured" (Cassam 2019, 126), we have revision responsibility for them, which is sufficient for blame when things go wrong.

Battaly (2016) also thinks that the provenances of bad traits do not matter in our vice attributions, but she goes further than Cassam. Battaly's personalist vice is entirely indifferent to agents' control (Battaly, 2015, 2016). For instance, she maintains that we blame the cruelty of Robert Harris or Hitler's youth even if we know that their unfortunate circumstances might have taken more prominent roles in producing their traits than self-cultivations. Battaly thinks such agents retain attributability responsibility, a responsibility in the sense that a relevant trait reflects badly on the person, which is sufficient for the trait to be vicious.

Cassam's theory and Battaly's theory extend the terrain of vices, including the traits that are acquired in a variety of ways. But how far the terrain should be

extended? Can the terrain really include the traits without responsibility in the sense of accountability (or backwards-looking responsibility) or malleability (or forward-looking responsibility)? It seems too much of a stretch to call such traits "vices" as we currently use the term; calling them "vices" requires drastic changes in our practices of charging other people. The more we are required to change our concept and practice concerning vice, the more the project of conceptualizing vice becomes ameliorative rather than descriptive (Haslanger, 2012).

If the ability to change our traits (i.e. revision responsibility) is not necessary for the traits to be vicious, then our "vice" talk may be so shallow and insubstantial that it is difficult to see the point of talking about "vices" in the first place. This leads to another requirement for vice theorists: the concept of "vice" is useful when attributing vices to someone can fulfill the "ameliorative task" (Kidd, 2016) of making the world a better place by either improving the intellectual character of the person or by improving the epistemic dynamic of the society in which the person is embedded.

Let us think again about the invisible-hand model. Can the group that emerged from the invisible-hand mechanism be responsible in any interesting sense? First, it is clear that invisible-hand groups are not responsible in the sense of accountability. The invisible-hand mechanism is, by definition, indifferent to how the relevant trait is formed. There is no interesting sense in which the emerging group trait out of the invisible-hand mechanism can be described as "chosen" or "prefered". Second, invisible-hand groups might not be revision responsible either. A problem is that invisible-hand groups may appear or disappear very swiftly, lacking unity as a group for an extended time. For a group to be vicious in the sense of revision responsible, the group needs minimal diachronic unity such that it makes sense to talk about reconfiguring the group trait for the future, or it makes sense to talk about the ameliorative task of improving the intellectual character of the group.

This is, of course, not to deny Holroyd's (2020) suggestion of implicit gender bias as a collective epistemic vice. A group displaying systematic implicit gender bias may actually have something that unifies the group such that the bias can be regarded as a collective epistemic vice of the group. In particular, they may be united by some group identification (e.g. as men, academic philosophers, or people working in BBC, etc.). In fact, the actual examples of the group in Holroyd's discussion are often groups with some group identities, such as a group of academic philosophers. Thus, implicit gender bias may be rightly called a collective epistemic vice from the viewpoint of GIA.

6 Conclusion

We started with a tension between the theoretical and applied aims of collective vice epistemology. On the theoretical side, Fricker's JCA is philosophically sophisticated, and it is consistent with our everyday ascription of epistemic virtues/vices to established groups. On the applied side, in contrast, JCA is not easily applicable to loose groups in real-life cases such as a loose group in an echo chamber. As a possible solution, this paper explored an alternative account of collective epistemic vices that is applicable not only to established groups but also to loose groups. As a good can-

didate for such an account, we presented GIA, which attributes collective epistemic vices to the groups that are constituted by group identification.

As we already noted in the end of Sect. 4, a limitation of our argument is that there are some open empirical issues. Our hypothesis is that people in echo chambers are motivated by group identification such that, according to GIA, they are collectively vicious. But the truth of this hypothesis depends on open empirical issues as to how exactly echo chambers operate and how people think and behave in them. Relevant empirical disciplines (psychology, sociology, etc.) do not offer conclusive evidence on these issues at this moment.

Another limitation is that we only focused on echo chamber as a case study, and we did not discuss other possible cases of vicious loose groups, such as the case of implicit bias (Holroyd, 2020) or the case of (loose) group polarization (Broncano-Berrocal & Carter, 2021). Although we suspect that GIA is applicable not only to the case of echo chamber but also to a wide range of cases of vicious loose groups, we admit that careful philosophical discussions and informative empirical data are necessary.

Declarations

Competing interests The authors have no competing interests to declare that are relevant to the content of this article.

References

Anderson, E. (2021). Epistemic bubbles and authoritarian politics. In E. Edenberg, & M. Hannon (Eds.), *Political Epistemology* (pp. 11–30). Oxford University Press.

Avnur, Y. (2020). What's wrong with the online echo chamber: A motivated reasoning account. *Journal of Applied Philosophy*, *37*(4), 578–593.

Bacharach, M. (2006). In N. Gold, & R. Sugden (Eds.), *Beyond Individual Choice: Teams and Frames in Game Theory*. Princeton University Press.

Battaly, H. (2015). Epistemic virtue and vice: Reliabilism, responsibilism, and personalism. In C. Mi, M. Slote, & E. Sosa (Eds.), *Moral and Intellectual Virtues in Western and Chinese Philosophy* (pp. 109–130). Routledge.

Battaly, H. (2016). Developing virtue and rehabilitating vice: Worries about self-cultivation and self-reform. *Journal of Moral Education*, *45*(2), 207–222.

Battaly, H. (2022). Solidarity: Virtue or vice? In M. Alfano, C. Klein, & J. de Ridder (Eds.), *Social Virtue Epistemology* (pp. 303–324). Routledge.

Begby, E. (2022). From belief polarization to echo chambers: A rationalizing account. *Episteme*. https://doi.org/10.1017/epi.2022.14

Bermúdez, J. L. (2022). Rational framing effects: A multidisciplinary case. *Behavioral and Brain Sciences*. https://doi.org/10.1017/S0140525X2200005X

Bratman, M. E. (2009). Modest sociality and the distinctiveness of intention. *Philosophical Studies*, *144*(1), 149–165.

Brewer, M. B. (1991). The social self: On being the same and different at the same time. *Personality and Social Psychology Bulletin*, *17*(5), 475–482.

Brewer, M. B., & Gardner, W. (1996). Who is this "We"? Levels of collective identity and self representations. *Journal of Personality and Social Psychology*, *71*(1), 83–93.

Broncano-Berrocal, F., & Carter, J. A. (2021). *The Philosophy of Group Polarization: Epistemology, Metaphysics, Psychology*. Routledge.

Byerly, T. R., & Byerly, M. (2016). Collective virtue. *The Journal of Value Inquiry*, *50*(1), 33–50.

Cassam, Q. (2019). *Vices of the Mind: From the Intellectual to the Political*. Oxford University Press.

Cassam, Q. (2021a). Misunderstanding vaccine hesitancy: A case study in epistemic injustice. *Educational Philosophy and Theory*, 1–15. https://doi.org/10.1080/00131857.2021.2006055

Cassam, Q. (2021b). Epistemic vices, ideologies, and false consciousness. In M. Hannon, & J. de Ridder (Eds.), *The Routledge Handbook of Political Epistemology* (pp. 301–311). Routledge.

Crerar, C. (2018). Motivational approaches to intellectual vice. *Australasian Journal of Philosophy*, *96*(4), 753–766.

de Ridder, J. (2022). Three models for collective intellectual virtues. In M. Alfano, C. Klein, & J. de Ridder (Eds.), *Social Virtue Epistemology* (pp. 367–385). Routledge.

de Rooij, B., & de Bruin, B. (2022). Real-life collective epistemic virtue and vice. In M. Alfano, C. Klein, & J. de Ridder (Eds.), *Social Virtue Epistemology* (pp. 396–414). Routledge.

Diehl, M. (1990). The minimal group paradigm: Theoretical explanations and empirical findings. *European Review of Social Psychology*, *1*(1), 263–292.

Fricker, M. (2010). Can there be institutional virtues? In T. S. Gendler, & J. Hawthorne (Eds.), *Oxford Studies in Epistemology Vol* (3 vol., pp. 235–253). Oxford University Press.

Fricker, M. (2020). Institutional epistemic vices: The case of inferential inertia. In I. J. Kidd, H. Battaly, & Q. Cassam (Eds.), *Vice Epistemology* (pp. 89–107). Routledge.

Gilbert, M. (1987). Modelling collective belief. *Synthese*, *73*(1), 185–204.

Gilbert, M. (1989). *On Social Facts*. Princeton University Press.

Gilbert, M. (2013). *Joint Commitment: How We Make the Social World*. Oxford University Press.

Haslanger, S. (2012). *Resisting Reality: Social Construction and Social Critique*. Oxford University Press.

Holroyd, J. (2020). Implicit bias and epistemic vice. In I. J. Kidd, H. Battaly, & Q. Cassam (Eds.), *Vice Epistemology* (pp. 126–147). Routledge.

Jamieson, K. H., & Cappella, J. N. (2008). *Echo Chamber: Rush Limbaugh and the Conservative Media Establishment*. Oxford University Press.

Kidd, I. J. (2016). Charging others with epistemic vice. *The Monist*, *99*(2), 181–197.

Kidd, I. J. (2021). Epistemic corruption and political institutions. In M. Hannon, & J. de Ridder (Eds.), *The Routledge Handbook of Political Epistemology* (pp. 347–358). Routledge.

Korsgaard, C. M. (1996). *The Sources of Normativity*. Cambridge University Press.

Lahroodi, R. (2018). Virtue epistemology and collective epistemology. In H. Battaly (Ed.), *The Routledge Handbook of Virtue Epistemology* (pp. 407–419). Routledge.

Levy, N. (2022a). *Bad Beliefs: Why They Happen to Good People*. Oxford University Press.

Levy, N. (2022b). Echoes of covid misinformation. *Philosophical Psychology*. https://doi.org/10.1080/09515089.2021.2009452

Medina, J. (2021). Vices of the privileged and virtues of the oppressed in epistemic group dynamics. In M. Hannon, & J. de Ridder (Eds.), *The Routledge Handbook of Political Epistemology* (pp. 336–346). Routledge.

Millikan, R. G. (1995). Pushmi-pullyu representations. *Philosophical Perspectives*, *9*, 185–200.

Millikan, R. G. (2004). *Varieties of Meaning: The 2002 Jean Nicod Lectures*. MIT press.

Miyazono, K., & Inarimori, K. (2021). Empathy, altruism, and group identification. *Frontiers in Psychology*. https://doi.org/10.3389/fpsyg.2021.749315. 5775.

Nguyen, C. T. (2020). Echo chambers and epistemic bubbles. *Episteme*, *17*(2), 141–161.

Pacherie, E. (2013). Intentional joint agency: Shared intention lite. *Synthese*, *190*(10), 1817–1839.

Payne, B. K., Vuletich, H. A., & Lundberg, K. B. (2017). The bias of crowds: How implicit bias bridges personal and systemic prejudice. *Psychological Inquiry*, *28*(4), 233–248.

Pickard, H. (2021). Addiction and the self. *Noûs*, *55*(4), 737–761.

Quine, W. V., & Carnap, R. (1991). *Dear Carnap, Dear Van: The Quine-Carnap Correspondence and Related Work*. University of California Press.

Salice, A., & Miyazono, K. (2020). Being one of us. Group identification, joint actions, and collective intentionality. *Philosophical Psychology*, *33*(1), 42–63.

Smart, P. R. (2018). Mandevillian intelligence. *Synthese*, *195*(9), 4169–4200.

Tajfel, H., Billig, M. G., Bundy, R. P., & Flament, C. (1971). Social categorization and intergroup behaviour. *European Journal of Social Psychology*, *1*(2), 149–178.

Turner, J. C. (1982). Towards a cognitive redefinition of the social group. In H. Tajfel (Ed.), *Social Identity and Intergroup Relations* (pp. 15–40). Cambridge University Press.

Zagzebski, L. T. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge University Press.