

EPISTEMIC DEPENDENCE, COGNITIVE IRRATIONALITY, AND EPISTEMIC CONFLICTS OF INTERESTS: WHY THERE IS A NEED FOR SOCIAL EPISTEMIC NORMS

Basil MÜLLER

ABSTRACT: When an agent A depends on an agent B to promote one of A's epistemic goals, this will often involve B's forming and sharing of true beliefs. However, as is well documented in research on cognitive irrationality, agents are disposed to form and share false-but-useful beliefs in a lot of circumstances. The dependence relation is thus at risk of becoming negative: A might adopt false beliefs from B and thus be unable to promote their epistemic goal. I propose that we can employ the notion of an epistemic conflict of interest [ECOI] to capture the kinds of problems that epistemically interdependent agents face. Much like familiar cases of conflict of interests—e.g., related to government officials—in ECOI an agent is subject to a normatively primary interest—roughly to form and share true beliefs—that stands in conflict with normatively secondary interests. I focus on secondary interests documented in the aforementioned research on cognitive irrationality. The resulting framework addresses an explanatory gap in the literature on social epistemic norms by making explicit why there's a need for these norms to regulate our epistemic lives. Lastly, I show how the ECOI-framework furthermore allows us to make sense of and amend norm regulation failures.

KEYWORDS: cognitive irrationality, epistemic dependence, social norms, epistemic norms, conflict of interests, epistemic conflicts of interest

1. Introduction

Much recent work in social epistemology has focused on instances of agent-based dependence, where an agent A depends on an agent B to promote one of A's epistemic goals (Goldberg 2011; Hardwig 1985; Pritchard 2015; for a recent overview see Broncano-Berrocal & Vega-Encabo 2020). Rather often, this involves B's forming and sharing of true, rational, or justified beliefs about topics of relevance to A.

At the same time, work on cognitive irrationality documents how agents form their beliefs not just with an eye to truth, but to what's useful or beneficial to them

(Gigerenzer & Selten 2002; Kunda 1990; for a short overview, see Williams 2021c). For example, research on identity protective cognition [IPC] describes the tendency of individuals to sample and process information with the aim of protecting or enabling their status as a member of a desirable social group.

Considered in combination, a systematic problem arises: A depends on B to form and share true beliefs (say about the climate crisis), yet B is inclined to form and share false-but-useful beliefs (say to belong to a climate-skeptic social group). A might thus acquire false beliefs which would prohibit A from promoting their epistemic aim—the dependence relation is at risk of becoming negative.

The main aim of this paper then is to provide a framework to capture and systematize these and related cases. To that end, I propose that we can employ the notion of an epistemic conflict of interests [ECOI]. ECOI are conflicts of interest over how to form and share beliefs. As in standard cases of conflicts of interests, in ECOI a normatively primary interest (e.g., to form and share true beliefs about topics of relevance) stands in conflict with normatively secondary interests (e.g., to belong to a social group).

The ECOI-framework not only provides a way to capture, categorize, and relate different kinds of risks epistemically interdependent agents are vulnerable to (see e.g. Broncano-Berrocal & Vega-Encabo 2020, sec. 6.4; Grasswick 2004), it also is of relevance to the literature on social epistemic norms. Several social epistemologists engaged in a variety of debates rely on or argue for the claim that (at least some) epistemic norms are social norms (Abbate 2021; Faulkner 2011; Goldberg 2020b; Graham 2012; 2015; Greco 2020; Henderson & Graham 2019; Simion 2021). However, it is standardly acknowledged that social norms emerge to regulate cooperation, where there's an incentive for individuals to defect, cheat, or free-ride (see e.g. Bicchieri 2005; Fehr, Fischbacher, & Gächter 2002; Henrich & Muthukrishna 2021)—that is once it is in an individual's, but not their groups' interest to behave in non-conforming ways. When it comes to social *epistemic* norms specifically, we find that a general conceptual framework to capture and classify the different ways in which individuals might obstruct epistemic cooperation is lacking. It is thus unclear why there's a need for social epistemic norms.¹ I will argue that this lack of understanding also has implications for our chances to employ social epistemic norms effectively, since incentivizing individuals away from a particular epistemic behavior requires us to understand that behavior and the interests it promotes in the first place.

Summarizing, the aims of this paper are, first, to provide a framework to capture the problems epistemically interdependent agents are exposed to and,

¹ But see Henderson (2020) for an important distinction that I'll revisit later on.

second, to consider how social epistemic norms regulate our epistemic behavior in light of them. Thus understood, this paper differs from existing research in a few ways:

Although there is a lively literature in applied ethics on general conflicts of interest (e.g., Almassi 2017; Brody 2011; Rodwin 2018; Wiersma et al. 2020) outside of select professional contexts their epistemic dimension has only been given little attention.

The notion of epistemic conflicts of interests has been implicitly and explicitly acknowledged in empirical research on cognitive irrationality, much of which I will discuss below. See for example Kahan with regards to IPC (2012, 732): “[...] public divisions over climate change stem [...] from a distinctive conflict of interest: between the personal interest individuals have in forming beliefs in line with those held by others [...] and the collective one they all share in making use of the best available science to promote common welfare.” However, there’s been no comprehensive treatment of these cases in social epistemology up to date.

Lastly, the literature on the ethics of beliefs is inspired by and discusses cases similar to the ones mentioned above, but there are also a few important differences: This is a piece of descriptive, rather than normative epistemology. What I mean by this is that—although the account is compatible with lots of positions in these debates—I’m not arguing for what agents *ought* to do in cases where epistemic and practical reasons clash (see e.g. Rinard 2019; Schmidt 2021), or how we might go about comparing these reasons (see e.g. Meylan 2021), or what gives epistemic reasons their normative force (e.g. Steglich-Petersen 2018). Rather, I provide a framework that allows us to describe and understand i) why and how these cases are relevant to extant social groups and ii) the actual and at times imperfect regulatory solutions they seek to employ in light of them.

The resulting account of ECOI then not only furthers our understanding of these matters but in doing so also enhances the prospects for epistemic norms to regulate our epistemic lives more efficiently.

2. Epistemic Interdependence and Cognitive Irrationality: A Systematic Problem

This section serves to introduce the fundamental sort of problem that the ECOI-framework is supposed to capture. Here I give a brief but hopefully helpful overview of research on epistemic (inter-)dependence and cognitive irrationality, before drawing attention to the problem that arises once they are considered in combination.

2.1 Epistemic Interdependence

Humans are *ultrasocial* beings: We intensely learn from and cooperate with others to attain various goods we'd be unable to attain on our own. It is well documented that this creates large interdependence within social groups: We depend on others for the success of many of our actions and the welfare of our social groups. In turn, this interdependence has had profound and lasting impacts on human cognition and the organization of our social groups (see e.g. Henrich & Muthukrishna 2021; Sterelny 2012; 2021; Tomasello et al. 2012).

In present-day (social) epistemology it is widely accepted that this interdependence is not only practical — it also has an epistemic dimension. We depend on others, not just for the success of our individual or collective actions, but also to attain individual and collective epistemic goals (see amongst others Broncano-Berrocal & Vega-Encabo 2020; Goldberg 2011; Henderson 2020; Henderson & Graham 2017b; 2017a; Pritchard 2015).

To be a bit more specific, I am interested in what is sometimes called agent-based dependence, where an agent A depends on an agent B (or B's epistemic standings² and practices) to attain some epistemic goal.^{3,4} Such agent-based dependence can take various forms and minimally occurs in: i) testimony (Coady 1992; Goldberg 2011), ii) collaborative and coordinated epistemic activities and projects (Hallsson & Kappel 2020)—e.g. in science (Kitcher 1990; De Ridder 2014), iii) our more diffusely depending on others to monitor, police, or apply the necessary epistemic vigilance (Sperber et al. 2010) to the beliefs and belief-forming practices of others (Goldberg 2011), and finally in iv) us being depended on epistemic instruments and epistemically engineered environments (S. Goldberg 2020a).⁵

The above considerations make clear that human social groups are practically and epistemically interdependent. Positive instances of epistemic (inter-)dependence can of course be enormously beneficial for social groups, allowing for epistemic division of labor, specialization, and the accumulation of culture and knowledge (Goldberg 2011; Sterelny 2012; 2021). But negative epistemic dependence makes social groups vulnerable: They're at risk of being obstructed in

² i.e., whether their beliefs are justified, true, knowledgeable, etc.

³ I plan to stay neutral on the question of whether collectives or groups can count as genuine epistemic agents, though I confess to having sympathies for Lackey's view (Lackey 2020; 2015).

⁴ For other kinds of epistemic dependence and a more general overview of the field see e.g. Broncano-Berrocal & Vega-Encabo (2020) and the articles in the related special issue.

⁵ It remains controversial whether depending on epistemic instruments and epistemically engineered environments can be reduced to agent-based dependence (Goldberg 2020a; Broncano-Berrocal & Vega-Encabo 2020).

their epistemic functioning and, in turn, in their overall welfare by a variety of factors. It is to *one* of these factors that I now turn.

2.2 Cognitive Irrationality

Human social groups might be epistemically interdependent, yet at the same time, it is well documented in empirical and theoretical research on cognitive irrationality that humans form their beliefs not (just) with an eye to truth but to what is useful or beneficial to them (Kahneman 2011; Kunda 1990; Gigerenzer & Selten 2002). Here I will focus on two of the most prominently discussed explanations of cognitive irrationality⁶ referring to processing-costs and motivated cognition respectively (Williams 2021a; 2021c).

The literature on processing-costs seeks to explain cognitive irrationality by means of constraints of time, resources, and computational power on human cognition (Gigerenzer & Selten 2002; Kahneman 2003; Lieder & Griffiths 2020). It is because we are creatures faced with these constraints that we need to find cost-effective ways of forming beliefs. Thus, we often rely on what is called ‘fast-and-frugal’ heuristics, as investigated by e.g., Gigerenzer & Selten (2002). These heuristics are simple, task-specific decision strategies for solving judgment and decision tasks in the most effective way possible, given the constraints mentioned above.

Although these heuristics lead to systematic mistakes, it is widely held that it is more beneficial for the individual to rely on them, rather than a less mistake-prone but more resource-intensive way of forming beliefs—which in principle would be available to the agents and are used if the belief in question is of a certain personal or social importance (Mercier & Sperber 2017; Lieder & Griffiths 2020). So even though individuals behave *epistemically* irrationally by not employing a reliable or otherwise truth-conducive way of forming beliefs, they still behave ecologically rationally (Gigerenzer 2008): Given the environments they inhabit and the constraints they face, they apply a decision-strategy that most effectively achieves their interests.

Research on the remaining source of cognitive irrationality—motivated cognition and reasoning—maintains that our motivations—e.g., our desires, aims, wants, goals—can causally influence the ways in which we form beliefs (Bénabou & Tirole 2016; Kahan 2017; Kunda 1990; Sharot & Garrett 2016; Williams 2021a).

⁶ By ‘*cognitive* irrationality’ I simply mean information-processing that is “systematically biased away from the truth” (Williams 2021b, 7) (for whichever reason this might be). I don’t make any normative claims as to whether agents ought (not) to process information in this way—this, to my mind, would be a question of *epistemic* rationality or normativity. I say more on this in section 4.

Explanations invoking motivational influences can be found in both psychology (Kunda 1990) and the social sciences (Bénabou & Tirole 2016). The underlying picture is the same: Beliefs have both epistemic and non-epistemic effects. The former are the effects of a belief that are related to the truth of its content, the latter are the effects that are not related to the truth of its content (Williams 2021a). Individuals can be motivated to attain the non-epistemic effects of (not)⁷ having a belief. These range from, amongst others, preserving or bettering their self-image or self-esteem (McKay & Dennett 2009; Zimmermann 2020), to overcoming self-control problems (Bénabou & Tirole 2016; 2002), to achieving emotional regulation (McKay & Dennett 2009; Sweeny et al. 2010; Sharot & Garrett 2016) and/or to promoting a particular social outcome (Williams 2021a; Zimmermann 2020)—think of the cases of IPC, where individuals aimed to protect or enable their status as a member of some social group by forming beliefs about policy-relevant topics. IPC is thus thought to be responsible for disagreement regarding policy-relevant facts—e.g., about the climate crisis or vaccine safety (Kahan 2012; 2017; Kahan et al. 2012; Van Bavel & Pereira 2018).⁸ A variety of unconscious or sub-personal processes are employed to arrive at the desired belief. For example, individuals will selectively recall evidence relevant to their desired belief (Kunda 1990; Bénabou & Tirole 2016). Motivated reasoning is found in all sorts of individuals, ranging from lay-people to well-educated individuals (Kahan 2017) and CEOs (Malmendier & Tate 2005; 2008).

Summarizing, what the research on cognitive irrationality tells us is that we tend to form false-but-useful beliefs, useful either because the *process* of forming the respective belief was cost-effective or because (not) *having* the respective belief accords with our motivations.

2.3 A Systematic Problem

Consider the following case [FARAH&JASMINE]:

Farah is a student in Jasmine's class. As part of their school's curriculum, Jasmine is

⁷ Individuals cannot only be motivated to form a belief but can also be motivated to remain ignorant (see Williams 2021b on motivated ignorance).

⁸ It is controversial whether individuals' beliefs in these cases are the result of motivated reasoning rather than e.g. limited scientific literacy or prior beliefs about the topic at hand (Pennycook & Rand 2019; Tappin, Pennycook, & Rand 2021). Note that IPC only functions as an example in this instance. The more general point, that individuals can be motivated to form beliefs to attain specific social outcomes, can be made without relying on IPC (Williams 2021a). There is also some controversy about whether individuals do in fact have different *beliefs* (as opposed to other, non-doxastic attitudes) about policy-relevant facts. I say more on this in section 7.

Epistemic Dependence, Cognitive Irrationality, and Epistemic Conflicts of Interests

supposed to teach Farah and the other students about the climate crisis. However, next to being a teacher, Jasmine is also a proud member of a local political party that recently has become rather climate sceptic. Though Jasmine hasn't felt strongly about the topic before, she has found herself in agreement with her party's views. In addition, she fears that her reputation in the party would be tarnished, were she to teach what's on the curriculum. So, instead of teaching about the environmental effects of factory farming, Jasmine focuses on the kinds of climate sceptic arguments appreciated by her party. Farah, influenced by her teacher in her beliefs about the climate crisis, both accepts Jasmine's teachings and shares them with friends and family who now consider joining Jasmine's party.

FARAH&JASMINE is illustrative of the kinds of problems that epistemic agents face. On the one hand, they are epistemically interdependent—i.e., agents depend on others to engage in various epistemic practices, to form and share true, justified, and rational beliefs for the dependence relation to be positive, much like Farah depends on Jasmine for beliefs about the climate crisis. On the other hand, agents seem to be disposed to act and believe in ways that stand in conflict with others' needs, e.g., by forming and sharing false-but-useful beliefs—much like Jasmine shared her socially motivated and false beliefs about the climate crisis. As a result, Farah was unable to attain true, accurate, or justified beliefs about the matter. Because of the importance of the topic, this can have negative implications for the epistemic functioning and overall welfare, both of the epistemically dependent individual and—because of the epistemically interdependent nature of social groups, e.g., because Farah shared her newly acquired beliefs with friends and family—for the entire social group.

As mentioned, I propose that we can capture and engage with this problem by employing the notion of an epistemic conflict of interest. Most of us will be familiar with conflicts of interests [COI] in relation to politicians, government officials, company executives, researchers, or physicians. In COI, agents promote normatively secondary interests (e.g., monetary gain) instead of normatively primary interests (e.g., for physicians: treating the unwell in the best way possible). If COI are not regulated, they pose a risk for agents and their social group (see e.g. Almassi 2017; Brody 2011; Rodwin 2018; Wiersma et al. 2020).

Epistemic COI are COI concerning how to form and share beliefs. Here too agents—such as Jasmine—promote normatively secondary interests (e.g., forming false-but-useful beliefs to belong to a social group) instead of normatively primary interests (e.g., forming and sharing true beliefs). Since ECOI are a kind of conflict of interest [COI], it will be helpful to first, rather briefly, establish what COI are to then subsequently discuss ECOI.

3. What are Conflicts of Interests?

Let's start with a familiar example of a COI [LARRY]:

A physician called Larry receives money from a pharmaceutical company to dispense their medication. However, the medication in question is not the best available treatment for the patient.

In this case, Larry has a personal (monetary) interest to dispense the medication and another (conflicting) interest to give the patient the best available treatment. Larry cannot promote both interests at the same time. It is also clear that the latter interest is normatively speaking more significant than the former—Larry should prescribe the best medication—and that this latter interest is somehow related to Larry's being a physician.

LARRY matches a broadly accepted characterization of COI, which, roughly put, goes something like this: In COI an agent holds multiple interests, such that promoting one interest would go against promoting another interest, where the latter is normatively speaking more important than the former and due to a social role (Wiersma, Kerridge, & Lipworth 2018; Lipworth, Ghinea, & Kerridge 2019).

To make sense of this we can notice that in COI there are two conflicting interests, but one is of more normative weight. I will adapt Thompson's (1993) now widely used terminology to differentiate between the two competing interests: The normatively more significant interests are called *primary* interests. The normatively less significant interests are called *secondary* interests.⁹

The characterization also mentions that the primary interest should be due to a social role. To that end, we can differentiate between social and personal interests: *Personal* interests are the values, aims, goals, wants, or ends that agents pursue by means of actions and attitudes (Wiersma, Kerridge, & Lipworth 2018; Lipworth, Ghinea, & Kerridge 2019). In LARRY, the relevant personal interest is Larry's aim to attain more wealth. *Social* interests are the duties, obligations, or requirements that stem from social roles individuals are part of (Wiersma, Kerridge, & Lipworth 2018; Lipworth, Ghinea, & Kerridge 2019). Again, in LARRY, the relevant social interest is the requirement to allow for and contribute to the best possible treatment of the patient and that this interest is due to Larry's social role of being a physician.

For present purposes, we can understand social roles to be characteristic or socially expected patterns of behaviors and attitudes that are associated with individuals that inhabit specific positions within a social group. There is a point or

⁹ There might also be cases where both interests are of equal normative weight, but where we still cannot promote one without going against the other. I do not discuss such cases here—in part because it is unclear whether there's a need to regulate these kinds of cases.

purpose that corresponds to the respective patterns of behavior and attitudes. There is something these patterns of behaviors and attitudes are to achieve or attain—something agents *should* do given their social roles (Banton 1965; Biddle 1986; Haslanger 2012; Komarovskiy 1992).¹⁰ For example, the purpose of the social role of a physician is to treat the unwell, and doing so is the characteristic pattern of behavior that we come to expect from them. So, social roles, in part due to their point or purpose and the social expectations that are constitutive of them, require certain actions or attitudes from an individual. These kinds of requirements on an individual that inhabits a social role are meant by ‘social interest’ here.¹¹

COI, then, are conflicts between a primary social interest and secondary personal or social interests. This requirement is in place to capture the socio-normative relevance of COI: The promotion of a secondary interest over a primary social interest has potentially problematic consequences for the wellbeing or functioning of a social group and thus stands in need of being regulated by it. This is intuitive for well-known cases of COI involving physicians, politicians, researchers, or government officials. These are cases in which a secondary interest is being promoted (e.g., for LARRY: prescribing certain medication for monetary reward) instead of a primary interest related to the respective social role (e.g., for LARRY: to treat the unwell in the best way possible), where there are potentially problematic consequences for the social group (e.g., the health of group-members).

It is a consequence of this view that cases where the primary interest is personal rather than social do not qualify as a COI. This might strike some as counter-intuitive. But notice that paradigmatic cases of COI—COI involving physicians, politicians, researchers, or government officials—all revolve around social interests. We can thus see that the notion of COI is intimately tied to the conflict being relevant to a social group or its members. One might treat ‘COI’ as a technical term, meant to capture instances where primary social interests are at risk of not being promoted.

4. Epistemic Conflicts of Interests

This section serves to introduce epistemic COI. Here I first introduce the primary interest involved in ECOI, before focusing on secondary interests.

¹⁰ Though I’m mainly relying on Haslanger’s account, other accounts of social roles (e.g. Searle 1995; 2010) would work equally well for my purposes here.

¹¹ Strictly speaking, on such an understanding of social interests, they aren’t something that a subject has, but rather something that applies to them. For ease of use, I will continue the ‘looser’ formulation.

4.1 The Primary Epistemic Interest

The general case for ECOI starts with establishing what the competing interests are going to be. Importantly, given that ECOI are a kind of COI, at least the primary interest needs to be social in nature. I contend that the primary interest in ECOI is the requirement to reliably form and share true and avoid forming or sharing false beliefs (about a certain class of propositions that are of relevance to others). For ease of reference, I will call this the primary epistemic interest [PEI].

Here are a few initial remarks on this formulation:

- The PEI as it is formulated here is clearly an abstraction. What it would precisely mean for agents to reliably form and share true beliefs and avoid doing so for false beliefs (about a certain class of propositions that are of relevance to others) will depend on a variety of factors. To illustrate, James (James 1979) already discusses how believing truly and not believing falsely might require different epistemic behaviors from agents. Which behavior to display clearly depends, in part, on contextual factors and will thus not be discussed here.
- ‘Form’ in this formulation is shorthand for many different kinds of doxastic behavior and should minimally be understood to include the revision, regulation, suspension, and maintenance of beliefs (see e.g., Meylan 2013).
- I have opted for a folk-psychological notion of belief over more credence-based or Bayesian ones. However, this shouldn’t matter to the key aspects of ECOI.
- Furthermore, I here focus on ECOI that are about acquiring *true* whilst avoiding *false* beliefs. Though ECOI related to other epistemic notions such as coherence, justification, knowledge, understanding, or rationality are also of interest, truth offers a relatively simple case to start building up our understanding of ECOI.¹²

Here’s why I think that the PEI is a *social* interest: I take it that the PEI is a requirement that stems from the social role of being a Knower. Knowers are expected to reliably make public assertions (about a certain class of propositions that are of relevance to others) that others may take to be true (Abbate 2021; Congdon 2018)—and are thus required to form and share true beliefs.

Most often, agents inhabit the social role of being a knower indirectly, in virtue of inhabiting another social role (Abbate 2021). Social roles can differ with regards to how central the role of Knower is to them. For some social roles, forming and sharing true beliefs (about a certain class of propositions that are of relevance to others) simply is their main purpose. For example, an academic or a teacher is required to reliably form and share true beliefs about a certain class of propositions

¹² In the following, I will simply use ‘form and share true beliefs’ and thus omit the second Jamesian goal for reasons of succinctness.

(whatever it is they are researching/teaching) that is (more or less) relevant to their social group because doing so is the purpose of their social role. Being a Knower is the main thing they are required to do as part of their social roles. For others, forming and sharing true beliefs (about a certain class of propositions that are of relevance to others) is what agents are required to do *so as to be able to fulfill* their social role. There's a sense in which these agents are more indirectly required to be Knowers by their respective social roles. For example, a physician is primarily required to treat the unwell. But to do so, physicians must form and share true beliefs about all kinds of things. So, agents can be required to be knowers in more or less direct ways.

Recall as well that social roles are socially expected patterns of behaviors and attitudes. Recent work in (social) epistemology has acknowledged that we do indeed have expectations regarding others' epistemic behavior and react negatively if these expectations aren't fulfilled. We (often legitimately) *expect* all kinds of people inhabiting social roles (e.g. physicians, but also neighbors, teachers, parents, teammates, citizens, cashiers, and friends) to form and share true beliefs about lots of things (Goldberg 2018). If they do not do so, we blame them and hold them responsible (Boult 2021; Meylan 2017; Tollefsen 2017). These, and many other social roles, are "socially approved sources of information," that allow for an epistemic division of labor in epistemically interdependent communities (Goldberg 2011; Greco 2020). It is in this sense that the PEI is a social interest that is due to the social role of a Knower.

We are now in a better position to appreciate why the PEI is formulated in the way it is—why the PEI requires agents to *reliably form and share* true beliefs (*about a certain class of propositions that are of relevance to others*).

- I. The PEI requires agents to *reliably* form and share true beliefs because most cooperation and dependence in social groups do not take place in the form of one-shot interactions. Much rather, cooperation is recurrent and often long-lasting (Henrich & Muthukrishna 2021) and so are the kinds of dependencies listed above. This is why randomly (or otherwise unreliably) formed true beliefs will not suffice.
- II. The PEI requires agents to reliably form and share true beliefs that are of *relevance to others* in the sense that 'relevant' beliefs here simply is intended to mean the sorts of beliefs in relation to which there is epistemic (inter)dependence. So agents are required to form true beliefs if others depend on them to do so.¹³

¹³ It is common—if not uncontroversial (Steglich-Petersen 2018)—to think that one is required to form true beliefs even in cases of trivial truths, i.e. truths that are of no interest to individuals or

It is worth noting that the scope of *relevant* beliefs might be larger than initially assumed: First, because of the large degree of specialization, epistemic (inter-) dependence is rampant in present-day social groups (Millgram 2015). Second, the scope of relevant beliefs does not only contain the beliefs agents actually depend on at a given point in time but also the beliefs and associated belief-forming capacities that agents might come to depend on in the future. In addition, it is simply difficult to *estimate* which beliefs will be of relevance to others, so airing on the side of caution is a worthwhile policy (Grimm 2009).

- III. The PEI requires agents to reliably *form and share* true beliefs about a certain class of propositions that are of relevance to others simply because forming false beliefs will lead the agent to share false beliefs—which, in turn, will often lead to others acquiring false beliefs.
- IV. The PEI requires agents to reliably form and share true beliefs *about a certain class of propositions* that are of relevance to others because we depend on different agents for different beliefs—people are Knowers about different topics. This is how a division of epistemic labor is achieved. In addition, the class of propositions individuals are required to form and share true beliefs about likely not only includes specialized beliefs but also beliefs that are either common knowledge or could easily have been known by an individual.

Lastly, let me say a few words on the PEI being a *primary* interest:

1. Let me first note that I aim to remain non-committal with regards to the question of what makes an interest primary, i.e., what makes an interest have more normative weight than another. Relatedly, the framework is compatible with but not committed to the PEI's normative importance or overall legitimacy being due to or grounded in the social role of being a Knower (Abbate 2021) or the expectations we have of others (Goldberg 2018). All I've suggested is that it's a descriptive fact that in many social groups there exists the social role of a Knower that requires certain things of agents. I'm thereby not committed to the normative claim that the legitimacy or normative importance of that requirement is also due to that particular social role (or the related expectations). I take it to be a strength of the ECOI-framework that it is compatible with a variety of meta-epistemological theories.
2. Because of the above, I do not claim that the PEI is a primary interest in any particular case. What I do claim is that the PEI *is* a primary interest in a substantial amount of cases. I take it that this is both intuitive and supported by the fact that social groups—by means of social norms and practices—are organized so as to ensure that the PEI is being complied with. It is because sharing

social groups (Kelly 2003). I do not take a position on this matter here. But note that the ECOI-framework is compatible with there being cases where individuals are required to form true beliefs about trivial truths (or other propositions of no interest to social groups). Because of the lacking social interest, these cases will simply not count as an ECOI.

Epistemic Dependence, Cognitive Irrationality, and Epistemic Conflicts of Interests

and forming true beliefs often (but certainly not always) is of more importance than the promotion of an agent's other interests that these norms and practices persist (see sct. 6 for more on this).

4.2 An Initial Taxonomy of Epistemic Conflicts of Interests

Section 4.1 established the primary epistemic interest at play in ECOI. To complete the picture, we must now turn to the secondary interests, the promotion of which would hinder the fulfilment of the PEI. Since the PEI remains the same, we can use the different kinds of secondary interests to establish a taxonomy of different kinds of ECOI.¹⁴ What kinds of interests could be relevant here?

As noted in Section 3, secondary interests can themselves either be personal or social interests. The PEI is the requirement to reliably *form* and *share* true beliefs (about a certain class of propositions that are of relevance to others). It follows that there are at least four kinds of secondary interests that can compete with the PEI and that there are thus at least four different kinds of ECOI: i) personal and ii) social interests that concern the formation and iii) personal and iv) social interests that concern the sharing of true beliefs (about a certain class of propositions that are of relevance to others) (see figure 1). To illustrate, recall FARAH&JASMINE. On the proposed framework, Jasmine is entangled in an ECOI: On the one hand, the PEI applies to her—she's required to form and share true beliefs as part of being a teacher—but on the other hand, she has secondary interests—being part of a social group—that conflict with it. Farah, in turn, is at risk of acquiring false beliefs.

As was the case for the PEI, in treating these interests as secondary I do not attempt to make a judgement about any particular case. It might sometimes be perfectly justified to put one's personal or other social interests before the PEI. At times, doing so might actually *contribute* to the epistemic functioning of an agent (see e.g. Bortolotti 2020). What I am claiming is that there is a significant number of cases—in part the kinds of cases similar to FARAH&JASMINE—where this isn't so, where the kinds of interests discussed below are secondary. For these cases I provide a framework.

One might also worry that the promotion of doxastic or epistemic interests—be they primary or secondary—has an overly voluntaristic ring to it. Two points in response: First, talk of interests doesn't commit me to saying that we have direct control in pursuing these interests. Indirect control over belief-forming processes suffices to promote certain interests over others (e.g., Meylan 2017). Second, social epistemic norms change the incentive-structure so that it becomes more

¹⁴ ECOI₁ is different from ECOI₂ in that ECOI₁ contains secondary interests of the kind₁ and ECOI₂ contains secondary interests of the kind₂.

advantageous for agents to form and share beliefs in specific ways. This also doesn't require agents to be able to believe at will, but rather for their belief-forming mechanisms to be sensitive to social rewards and punishments (Williams 2021a).

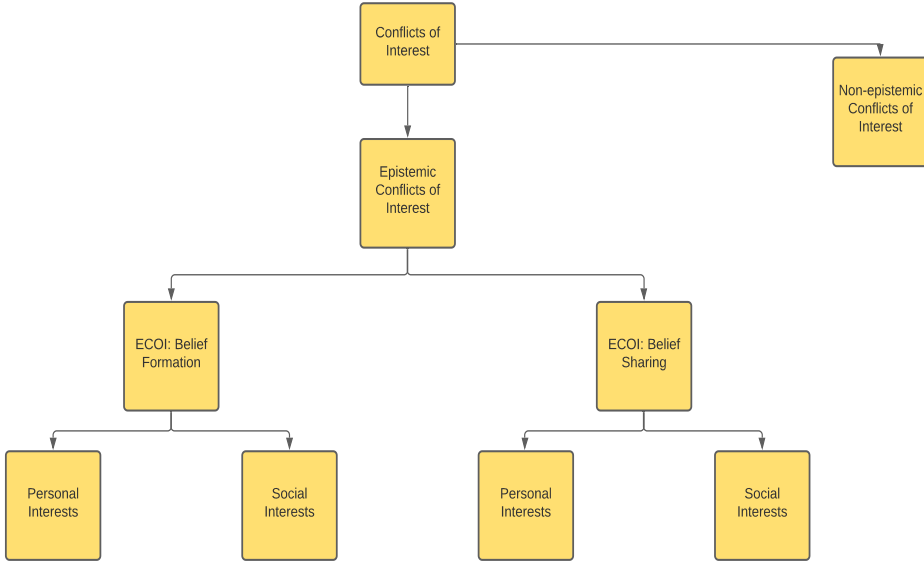


Figure 1: A taxonomy of different kinds of ECOI

I will focus on personal and social interests related to the *formation* of beliefs that might compete with the PEI. This is in part due to a lack of space. But more positively, drawing from this rich body of empirical work on cognitive irrationality allows for a relatively fine-grained and empirically adequate understanding of an important subset of ECOI.

But before turning to ECOI related to the *formation* of beliefs, let me acknowledge that there might be significant differences in what the PEI precisely amounts to in the two cases. For example, on Greco's (2016; 2020) view, forming and sharing beliefs have different social functions: Roughly put, the former serves to bring information into an epistemic community, whilst the latter functions to transmit it. To perform these functions effectively, agents might be required to behave in different ways.

Here are two examples for iii) personal and iv) social interests related to the *sharing* of beliefs that compete with the PEI to bring across the general idea, though both of these surely deserve more time and consideration. Personal interests related

to the sharing of beliefs that can conflict with the PEI are rather easy to come by: All of us lie from time to time. In some, if not most, cases we do so to promote some interest of ours. In contrast, whistleblower cases are instructive when it comes to social interests related to the sharing of beliefs: In such cases, there exist requirements that stem from the particular social role—often a particular profession—that forbid the agent from sharing information. Given the fact that individuals inhabit epistemically interdependent social groups, others will plausibly depend on them sharing true beliefs to a certain degree. Thus, ECOI arise.

This concludes the general overview of ECOI. In ECOI, the PEI—the requirement to reliably form and share true beliefs about a certain class of propositions—stands in conflict with both personal and social interests related to the formation and sharing of beliefs.¹⁵

5. Epistemic Conflicts of Interests Concerned with the Formation of Beliefs

This section serves to give a more fine-grained understanding of ECOI concerned with the *formation* of beliefs. I will rely on the two explanations of cognitive irrationality that I introduced in section 2.2. The different kinds of interests we find documented in this sort of research allow us to establish two ‘sub-kinds’ of ECOI and thus lets us expand on our initial taxonomy.

Before going into more detail, note again that I’m not committing myself to say that each case where supposedly cognitively irrational mechanisms are employed qualifies as an ECOI. This is, in part, because employing (some of these) mechanisms might at times be conducive to attaining good epistemic outcomes (see e.g., Hallsson & Kappel 2020; Mercier & Sperber 2017).

5.1 Process-ECOI

In Process-ECOI the interests of an individual might align with the PEI insofar as the individual, in principle, is interested in forming true beliefs about *p*. However, there is a conflict when it comes to how this interest is supposed to be pursued, i.e., how the belief is going to be formed.

¹⁵ ECOI so conceived bear resemblance to what is sometimes discussed as intellectual COI [ICOI] (Goldberg 2020; Wiersma, Kerridge, & Lipworth 2018) in applied ethics and adjacent areas. ICOI arise out of intellectual commitments to a particular research program or clinical practice that might bias the judgement of a physician or researcher (or similar). Though surely related, ECOI are a much broader notion in that they first do not only arise in relation to intellectual commitments but in relation to a variety of personal and social interests that compete with the PEI (see Sections 4.2-5.2). Second, ECOI are much broader in that they do not only arise in relation to a select few social roles but to a *multitude* of social roles (see Section 4.1).

Although such conflicts can manifest themselves in different ways, I take it that one of the most central instances is captured in processing-costs explanations of cognitive irrationality. Recall, these refer to constraints on time, resources, and computational power on human cognition (Gigerenzer & Selten 2002; Kahneman 2003; Lieder & Griffiths 2020). The general idea of these explanations is that because we are creatures faced with these constraints, we need to find *cost-effective* ways of forming beliefs.

There is a conflict regarding the costs associated with the formation of the respective belief. It arises because what is most cost-effective for the individual is not necessarily conducive to promoting the PEI in the manner required. Recall that individuals tend to, e.g., employ fast-and-frugal heuristics to cost-effectively form beliefs in such circumstances. Though useful (and certainly cost-effective), these lead to systematic mistakes. To *reliably* form true beliefs in different contexts—as required by the PEI—individuals would often need to form beliefs in ways that are not most cost-effective for them, e.g., by investing more time or computational power than they would like to.

So, Process-ECOI are about how individuals ought to form beliefs. Most centrally, they contain secondary interests related to the costs of forming beliefs.¹⁶

Before turning to Content-ECOI it should be noted that *social* interests might lead an agent to employ cost-effective but unreliable ways of belief-formation too. Consider a large clinic that is very demanding of its employees. Physicians have to see a large number of patients per day. This leads to them relying on fast-and-frugal heuristics to comply with the demands of their employer. At the same time, the quality of care suffers because of the lack of (epistemic) resources physicians are able to invest. Here physicians are entangled in a Process-ECOI, where a social interest conflicts with the PEI: As employees, they are (indirectly) required to form beliefs in unreliable ways by investing as little resources as possible. As physicians, they are required to do the opposite.

5.2 Content-ECOI

Whilst Process-ECOI are concerned with the costs of *forming* beliefs, Content-ECOI are concerned with (not) *having* specific beliefs. Motivated cognition explanations of cognitive irrationality refer to the fact that we are often motivated to (not) form specific beliefs and—because of this—employ unreliable ways of belief-formation.

¹⁶ So conceived, Process-ECOI bear similarities to what Henderson (2020, 290) calls ‘competing pursuit attractors.’ Both describe instances where individuals do not employ their best belief-forming capacities, in part because of the pursuit of other interests. Though note that Henderson doesn’t differentiate between social and personal interests.

The motivations to do so are related to beliefs' non-epistemic effects which range from self-regulation to the influence of one's beliefs on others.

In Content-ECOI individuals have interests to (not) form specific beliefs—i.e., beliefs with a specific content. Whilst individuals in Process-ECOI are at least in principle interested in forming true beliefs, in Content-ECOI individuals' interests to form specific beliefs are independent of the veracity of the beliefs. The interests are much rather related to the aforementioned non-epistemic effects of these beliefs.

The relevant interests can be both personal and social in nature. Content-ECOI with personal interests are rather straightforward: These are just the usual cases of individuals being motivated (e.g., to attain emotional regulation) to form specific beliefs. Content-ECOI with social interests are more complex. Cases of IPC are instructive here. In IPC an individual is motivated to (not) form a belief to protect or enable their status as a member of some desirable group. Though individuals in such cases have a personal interest, roughly to being part of their social group, frequently there will also be a corresponding social interest: A requirement for them to (not) have certain beliefs in order to be part of that social group. Having such beliefs thus becomes part of their social role—it is what is required and has come to be expected of them as part of their being a member of the group. Recall how in FARAH & JASMINE the latter does have a personal interest in being part of her respective political party. But it's also the case that to fulfil the social role of being a member of party X, Jasmine was required to have certain beliefs about the climate crisis. Promoting these kinds of interests might be conducive to being part of a social group, yet it will more often than not stand in competition with promoting the PEI: A belief formed for its non-epistemic effects will in most cases not be true.¹⁷

This concludes the discussion of Content-ECOI. I have argued that Content-ECOI arise due to personal and social interests to form specific beliefs irrespective of the veracity of the respective beliefs. More generally, this also concludes the discussion of the different kinds of ECOI related to the formation of beliefs. Content- and Process-ECOI are similar in that individuals employ unreliable or otherwise non-truth-conducive ways of belief-formation—this is how they infringe on the PEI. They are differentiated in that the secondary interests in Process-ECOI are related to the process of forming beliefs, whilst the secondary interests in Content-ECOI are concerned with (not) having specific beliefs.

¹⁷ So conceived, Content-ECOI are similar to what Henderson (2020, 290) calls 'content attractors.' Both describe cases where individuals have an interest (what Henderson calls a 'temptation') to form a belief because of its content and irrespective of its truth. Though note that Henderson doesn't differentiate between social and personal interests.




Figure 2: An expanded version of the taxonomy of ECOI—new additions in green

6. Epistemic Conflicts of Interests and Epistemic Social Norms

There are two related reasons that speak in favor of the ECOI-framework. First, it addresses an explanatory gap in the literature on social epistemic norms. Second, it is essential to properly understand instances of norm regulation failures.

With regards to the first point, a variety of social epistemologists engaged in different projects agree that there are social epistemic norms, which, roughly put, guide us in forming and sharing our beliefs (Abbate 2021; Dyke 2020; Faulkner 2011; Goldberg 2020b; Graham 2012; 2015; Greco 2020; Henderson & Graham 2019; Simion 2021). Such norms incentivize agents by means of sanctioning and reputation-mechanisms: We sanction individuals, in part by praising or blaming them for the beliefs they hold (Boult 2021). This also affects an individual’s epistemic reputation (Origgi 2012). In turn, we selectively associate with individuals who have a *sufficient* epistemic reputation—who we deem to be epistemically reliable (Henderson & Graham 2019, 429f).

However, it is standardly acknowledged that social norms (at least in the strict sense of the term mostly used in the literature) only emerge in cooperative contexts, when there are incentives for individuals to defect, cheat, or free-ride (see e.g., Bicchieri 2005; Fehr, Fischbacher, & Gächter 2002; Henrich & Muthukrishna 2021). Put more game-theoretically, social norms are seen to solve mixed-motive or cooperation games (e.g., a Prisoner's dilemma)—that is games where there is a good to be attained by cooperating but where the incentive-structure favors defecting over cooperating. Social norms solve such games by incentivizing agents to behave in the ways required for the good to be attained (Bicchieri 2005, chap. 1). However, if there are no incentives to defect then there's no need for social norms to regulate the kinds of social choice situations that cooperation games are supposed to model.

This general point also holds for social epistemic norms: If there are no incentives to defect from epistemic cooperation, there's no need for social epistemic norms.¹⁸ So far, only little has been said on how agents might defect from epistemic cooperation in the literature, most stipulate that such incentives likely do exist (but see Henderson 2020 for an exception). The ECOI framework addresses this explanatory gap by first, illuminating what agents are required to do or believe for epistemic cooperation to work out and by, second, offering a way to categorize different kinds of ways in which agents might be incentivized not to epistemically cooperate (see section 4.2 on secondary interests). It is because of the existence of ECOI that there's a need for social epistemic norms to regulate our epistemic behavior.

With regards to the second point, the ECOI-framework is essential to understand norm regulation failures and thus to securing successful epistemic cooperation within social groups. Of course, there might be many different things that go wrong in how social groups aim to regulate belief-formation. They might, for example, take the wrong interest to be primary. But there also exist distinct kinds of norm regulation failures, that is instances where norms fail to regulate their target behavior. I'll consider two here:

- 1) A social norm might fail to regulate its target behavior simply because the kind of incentivization it provides is insufficient. That is when agents—despite the existence of a social epistemic norm—still prefer to defect from epistemic cooperation. For example, in FARAH&JASMINE, Jasmine might come to believe as she does because the incentivization that is provided by social epistemic norms

¹⁸ In such instances, we might find there to be something like epistemic customs or conventions, that is social norms (in a broad sense of the term) that solve coordination games. I agree with Henderson (2020) that epistemic norms likely function as customs or conventions in some contexts and as social norms in others.

isn't so costly as to mitigate the benefits from believing in line with her social group. To make sense of and rectify this sort of situation, we need to acknowledge that Jasmine is part of an ECOI and adjust incentivization accordingly.

2) A social norm might also fail to regulate its target behavior because they address the wrong kind of ECOI and thus change the incentive-structure in the wrong way. For example, a social epistemic norm meant to regulate ECOI that concern how beliefs are shared will not be successful in regulating ECOI that are about the formation of beliefs—clearly, agents might only share beliefs they take to be true, yet all too often form false beliefs. Relatedly, social epistemic norms that seek to regulate Process-ECOI might not be successful in regulating Content-ECOI. Consider again FARAH&JASMINE: Incentivizing Jasmine to make use of her best reasoning-processes might not be successful if she's independently motivated to acquire specific beliefs. There's evidence that shows that people with a high degree of scientific literacy tend to employ these capacities to arrive at the desired conclusions (see e.g. Kahan et al. 2012). So Jasmine too might simply reason her way to the desired belief. This second point also shows why it is important to have a fine-grained understanding of different kinds of ECOI: The reason for this is simply that they might need to be regulated in different ways.

7. Conclusion

I have argued that the notion of an ECOI—as part of which the PEI stands in conflict with a variety of personal and social interests—can be used to capture a range of risks that epistemically interdependent agents are exposed to. In particular, the ECOI-framework allows us to understand how cognitive irrationality gives rise to a systematic problem for such agents and how social groups seek to address this problem by means of social epistemic norms.

There is of course much to further be addressed when it comes to attaining a fine-grained understanding of ECOI—both additional empirical *and* philosophical work is required.

On the empirical side, we can draw on research on IPC to elucidate a few important points. People skeptical of IPC-research question whether individuals do in fact not have different *beliefs* regarding policy-relevant facts. Rather, they claim that what explains an individual's assertions is what is often called 'expressive responding' (Bullock, Gerber, & Huber 2015) or 'motivated responding' (Khanna & Sood 2018)—i.e., individuals expressing their (non-factual) attitudes (e.g. of support or dislike) with regards to a certain topic. If individuals are lacking beliefs in these instances, then there will not be an ECOI related to the *formation* of beliefs. (Although, since their assertions are likely to be *interpreted* as expressions of beliefs (Hannon & de Ridder 2021, 159) they might be part of an ECOI that is related to the *sharing* of beliefs.) It is thus easy to see how more empirical work is required to more

accurately classify different kinds of ECOI. Of course, empirical work not considered here is also of importance for extending the ECOI-taxonomy.

On the philosophical side, two points stand out: First, ECOI related to sharing of beliefs have only been given little attention in this paper. This ought to be rectified. Second, ECOI related to cases of epistemic injustice (Fricker 2007), exploitation (Berenstein 2016), corruption (Kidd 2019), oppression (Sertler 2022), or even entitlement (Manne 2020) should be considered. In these and related cases there plausibly are both personal and social interests that hinder individuals from forming and/or sharing true beliefs. These are surely normatively secondary interests.

Lastly, non-epistemic COI are in part regulated by requiring the disclosure of potential COI (Giubilini & Savulescu 2020, 240f). What seems to be missing in present-day discourse is a requirement to disclose one's being part of an *epistemic* COI. However, doing so requires the capacity to recognize and acknowledge that one is, in fact, part of one. I hope that the conceptual framework provided in this paper is a first step in this direction.¹⁹

References

- Abbate, Cheryl. 2021. "On the Role of Knowers and Corresponding Epistemic Role Oughts." *Synthese* 199 (3): 9497–9522.
- Almassi, Ben. 2017. "Toxic Funding? Conflicts of Interest and Their Epistemological Significance." *Journal of Applied Philosophy* 34 (2): 206–20.
- Banton, M. 1965. *An Introduction to the Study of Role Relations*. Basic Books.
- Bénabou, Roland, and Jean Tirole. 2002. "Self-Confidence and Personal Motivation*." *The Quarterly Journal of Economics* 117 (3): 871–915. <https://doi.org/10.1162/003355302760193913>.
- . 2016. "Mindful Economics: The Production, Consumption, and Value of Beliefs." *Journal of Economic Perspectives* 30 (3): 141–64. <https://doi.org/10.1257/jep.30.3.141>.
- Berenstein, Nora. 2016. "Epistemic Exploitation." *Ergo, an Open Access Journal of Philosophy* 3: 569–90. <https://doi.org/10.3998/ergo.12405314.0003.022>.
- Bicchieri, Cristina. 2005. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge: Cambridge University Press.

¹⁹ **Acknowledgements:** I would like to express my gratitude to Rebekka Hufendiek, Anne Meylan, Ivan Gonzalez-Cabrera, Sebastian Schmidt, and the participants at both the Research Colloquium for Philosophy of Mind and Metaphysics at the University of Bern and the European Epistemology Network Meeting in Glasgow for their input and invaluable help in writing this paper. I would also like to thank the anonymous referees whose comments greatly benefitted this paper.

Basil Müller

- Biddle, Bruce J. 1986. "Recent Developments in Role Theory." *Annual Review of Sociology* 12 (1): 67–92.
- Bortolotti, Lisa. 2020. *The Epistemic Innocence of Irrational Beliefs*. New York: Oxford University Press.
- Boult, Cameron. 2021. "Epistemic Blame." *Philosophy Compass* 16 (8): e12762. <https://doi.org/10.1111/phc3.12762>.
- Brody, Howard. 2011. "Clarifying Conflict of Interest." *The American Journal of Bioethics* 11 (1): 7.
- Broncano-Berrocal, Fernando, and Jesús Vega-Encabo. 2020. "A Taxonomy of Types of Epistemic Dependence: Introduction to the Synthese Special Issue on Epistemic Dependence." *Synthese* 197 (7): 2745–63. <https://doi.org/10.1007/s11229-019-02233-6>.
- Bullock, John G., Alan S. Gerber, and Gregory A. Huber. 2015. "Partisan Bias in Factual Beliefs about Politics." *Quarterly Journal of Political Science* 10 (4): 519–78.
- Coady, Cecil Anthony John. 1992. *Testimony: A Philosophical Study*. Oxford: Clarendon Press ; Oxford University Press.
- Congdon, Matthew. 2018. "'Knower' as an Ethical Concept: From Epistemic Agency to Mutual Recognition." *Feminist Philosophy Quarterly* 4 (4). <https://doi.org/10.5206/fpq/2018.4.6228>.
- De Ridder, Jeroen. 2014. "Epistemic Dependence and Collective Scientific Knowledge." *Synthese* 191 (1): 37–53.
- Dyke, Michelle M. 2020. "Could Our Epistemic Reasons Be Collective Practical Reasons?" *Nous* 55 (4): 842–62. <https://doi.org/10.1111/nous.12335>.
- Faulkner, Paul. 2011. *Knowledge on Trust*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199589784.001.0001>.
- Fehr, Ernst, Urs Fischbacher, and Simon Gächter. 2002. "Strong Reciprocity, Human Cooperation, and the Enforcement of Social Norms." *Human Nature* 13 (1): 1–25. <https://doi.org/10.1007/s12110-002-1012-7>.
- Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.
- Gigerenzer, Gerd. 2008. *Rationality for Mortals: How People Cope with Uncertainty*. Oxford University Press.
- Gigerenzer, Gerd, and Reinhard Selten. 2002. *Bounded Rationality: The Adaptive Toolbox*. MIT press.
- Giubilini, Alberto, and Julian Savulescu. 2020. "Beyond Money: Conscientious Objection in Medicine as a Conflict of Interests." *Journal of Bioethical Inquiry* 17 (2): 229–43. <https://doi.org/10.1007/s11673-020-09976-9>.

- Goldberg, Daniel S. 2020. "Financial Conflicts of Interest Are of Higher Ethical Priority than 'Intellectual' Conflicts of Interest." *Journal of Bioethical Inquiry* 17 (2): 217–27. <https://doi.org/10.1007/s11673-020-09989-4>.
- Goldberg, Sanford. 2011. "The Division of Epistemic Labor." *Episteme* 8 (1): 112–25. <https://doi.org/10.3366/epi.2011.0010>.
- . 2018. *To the Best of Our Knowledge: Social Expectations and Epistemic Normativity*. Oxford, United Kingdom: Oxford University Press.
- . 2020a. "Epistemically Engineered Environments." *Synthese* 197 (7): 2783–2802. <https://doi.org/10.1007/s11229-017-1413-0>.
- . 2020b. "Social Epistemic Normativity: The Program." *Episteme* 17 (3): 364–83. <https://doi.org/10.1017/epi.2019.54>.
- Graham, Peter. 2012. "Testimony, Trust, and Social Norms." *Abstracta Special Issues* 6 (3): 92–116.
- . 2015. "Epistemic Normativity and Social Norms." In *Epistemic Evaluation: Purposeful Epistemology*, edited by John Greco and David Henderson, 247–73. Oxford University Press.
- Grasswick, Heidi E. 2004. "Individuals-in-Communities: The Search for a Feminist Model of Epistemic Subjects." *Hypatia* 19 (3): 85–120. <https://doi.org/10.1111/j.1527-2001.2004.tb01303.x>.
- Greco, John. 2016. "WHAT IS TRANSMISSION*?" *Episteme* 13 (4): 481–98. <https://doi.org/10.1017/epi.2016.25>.
- . 2020. *The Transmission of Knowledge*. Cambridge University Press.
- Grimm, Steven. 2009. "Epistemic Normativity." In *Epistemic Value*, edited by A. Haddock, A. Millar, and D. Pritchard, 243–64. Oxford: Oxford University Press.
- Hallsson, Björn G., and Klemens Kappel. 2020. "Disagreement and the Division of Epistemic Labor." *Synthese* 197 (7): 2823–47. <https://doi.org/10.1007/s11229-018-1788-6>.
- Hardwig, John. 1985. "Epistemic Dependence." *Journal of Philosophy* 82 (7): 335–49. <https://doi.org/jphil198582747>.
- Haslanger, Sally. 2012. *Resisting Reality: Social Construction and Social Critique*. New York: Oxford University Press.
- Henderson, David. 2020. "Are Epistemic Norms Fundamentally Social Norms?" *Episteme* 17 (3): 281–300.
- Henderson, David, and Peter Graham. 2017a. "A Refined Account of the 'Epistemic Game': Epistemic Norms, Temptations, and Epistemic Cooperation." *American Philosophical Quarterly* 54 (4): 383–95.

- . 2017b. "Epistemic Norms and the 'Epistemic Game' They Regulate: The Basic Structured Epistemic Costs and Benefits." *American Philosophical Quarterly*, 367–82.
- Henderson, David, and Peter J. Graham. 2019. "Epistemic Norms as Social Norms." In *The Routledge Handbook of Social Epistemology*, edited by Miranda Fricker, Peter J. Graham, David Henderson, and Nikolaj J. L. L. Pedersen, 425–36. Routledge. <https://doi.org/10.4324/9781315717937-42>.
- Henrich, Joseph, and Michael Muthukrishna. 2021. "The Origins and Psychology of Human Cooperation." *Annual Review of Psychology* 72 (1): 207–40. <https://doi.org/10.1146/annurev-psych-081920-042106>.
- James, William. 1979. *The Will to Believe and Other Essays in Popular Philosophy*. Harvard University Press.
- Kahan, Dan M. 2012. "Ideology, Motivated Reasoning, and Cognitive Reflection: An Experimental Study." *Judgment and Decision Making* 8: 407–24.
- . 2017. "Misconceptions, Misinformation, and the Logic of Identity-Protective Cognition." *Cultural Cognition Project Working Paper Series*. Yale Law School, Public Law Research Paper 605. SSRN. <https://doi.org/10.2139/ssrn.2973067>.
- Kahan, Dan M., Ellen Peters, Maggie Wittlin, Paul Slovic, Lisa Larrimore Ouellette, Donald Braman, and Gregory Mandel. 2012. "The Polarizing Impact of Science Literacy and Numeracy on Perceived Climate Change Risks." *Nature Climate Change* 2 (10): 732–35. <https://doi.org/10.1038/nclimate1547>.
- Kahneman, Daniel. 2003. "Maps of Bounded Rationality: Psychology for Behavioral Economics." *American Economic Review* 93 (5): 1449–75.
- . 2011. *Thinking, Fast and Slow*. Macmillan.
- Kelly, Thomas. 2003. "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 66 (3): 612–40. <https://doi.org/10.1111/j.1933-1592.2003.tb00281.x>.
- Khanna, Kabir and Gaurav Sood. 2018. "Motivated Responding in Studies of Factual Learning." *Political Behavior* 40 (1): 79–101. <https://doi.org/10.1007/s11109-017-9395-7>.
- Kidd, Ian James. 2019. "Epistemic Corruption and Education." *Episteme* 16 (2): 220–35. <https://doi.org/10.1017/epi.2018.3>.
- Kitcher, Philip. 1990. "The Division of Cognitive Labor." *The Journal of Philosophy* 87 (1): 5–22.
- Komarovsky, Mirra. 1992. "The Concept of Social Role Revisited." *Gender & Society* 6 (2): 301–13.

- Kunda, Ziva. 1990. "The Case for Motivated Reasoning." *Psychological Bulletin* 108 (3): 480.
- Lackey, Jennifer. 2015. "A Deflationary Account of Group Testimony." In *Essays in Collective Epistemology*, 64–94. Oxford University Press.
- . 2020. *The Epistemology of Groups*. Oxford, New York: Oxford University Press.
- Lieder, Falk, and Thomas L. Griffiths. 2020. "Resource-Rational Analysis: Understanding Human Cognition as the Optimal Use of Limited Computational Resources." *Behavioral and Brain Sciences* 43.
- Lipworth, Wendy, Narcyz Ghinea, and Ian Kerridge. 2019. "Clarifying the Relationship Between Serious Ethical Violations and Conflicts of Interest." *The American Journal of Bioethics* 19 (1): 48–50.
- Malmendier, Ulrike, and Geoffrey Tate. 2005. "Does Overconfidence Affect Corporate Investment? CEO Overconfidence Measures Revisited." *European Financial Management* 11 (5): 649–59. <https://doi.org/10.1111/j.1354-7798.2005.00302.x>.
- . 2008. "Who Makes Acquisitions? CEO Overconfidence and the Market's Reaction." *Journal of Financial Economics* 89 (1): 20–43. <https://doi.org/10.1016/j.jfineco.2007.07.002>.
- Manne, Kate. 2020. *Entitled: How Male Privilege Hurts Women*. Penguin Random House.
- McKay, R. T., and Daniel Dennett. 2009. "The Evolution of Misbelief." *Behavioral and Brain Sciences* 32 (6): 493–510.
- Mercier, Hugo, and Dan Sperber. 2017. *The Enigma of Reason*. Harvard University Press.
- Meylan, Anne. 2013. *Foundations of an Ethics of Belief*. Walter de Gruyter.
- . 2017. "The Consequential Conception of Doxastic Responsibility." *Theoria* 83 (1): 4–28. <https://doi.org/10.1111/theo.12110>.
- . 2021. "Doxastic Divergence and the Problem of Comparability. Pragmatism Defended Further." *Philosophy and Phenomenological Research* 103 (1): 199–216. <https://doi.org/10.1111/phpr.12709>.
- Millgram, Elijah. 2015. *The Great Endarkenment: Philosophy for an Age of Hyperspecialization*. Oxford University Press.
- Origgi, Gloria. 2012. "A Social Epistemology of Reputation." *Social Epistemology* 26 (3–4): 399–418. <https://doi.org/10.1080/02691728.2012.727193>.
- Pennycook, Gordon, and David G. Rand. 2019. "Lazy, Not Biased: Susceptibility to Partisan Fake News Is Better Explained by Lack of Reasoning than by

- Motivated Reasoning." *Cognition* 188: 39–50. <https://doi.org/10.1016/j.cognition.2018.06.011>.
- Pritchard, Duncan. 2015. "Epistemic Dependence." *Philosophical Perspectives* 29: 305–24.
- Rinard, Susanna. 2019. "Believing for Practical Reasons." *Noûs* 53 (4): 763–84. <https://doi.org/10.1111/nous.12253>.
- Rodwin, Marc A. 2018. "Attempts to Redefine Conflicts of Interest." *Accountability in Research* 25 (2): 67–78. <https://doi.org/10.1080/08989621.2017.1405728>.
- Schmidt, Sebastian. 2021. "Epistemic Blame and the Normativity of Evidence." *Erkenntnis*, 1–24.
- Searle, John. 1995. *The Construction Of Social Reality*. New York: Free Press.
- . 2010. *Making the Social World: The Structure of Human Civilization*. Oxford University Press.
- Sertler, Ezgi. 2022. "Epistemic Dependence and Oppression: A Telling Relationship." *Episteme* 19 (3): 394–408. <https://doi.org/10.1017/epi.2020.34>.
- Sharot, Tali, and Neil Garrett. 2016. "Forming Beliefs: Why Valence Matters." *Trends in Cognitive Sciences* 20 (1): 25–33.
- Simion, Mona. 2021. "Testimonial Contractarianism: A Knowledge-First Social Epistemology." *Noûs* 55 (4): 891–916. <https://doi.org/10.1111/nous.12337>.
- Sperber, Dan, Fabrice Clément, Christophe Heintz, Olivier Mascaro, Hugo Mercier, Gloria Origi, and Deirdre Wilson. 2010. "Epistemic Vigilance." *Mind & Language* 25 (4): 359–93.
- Steglich-Petersen, Asbjørn. 2018. "Epistemic Instrumentalism, Permissibility, and Reasons for Belief." In *Normativity: Epistemic and Practical.*, edited by Conor McHugh, Jonathan Way, and Daniel Whiting. Oxford University Press. <https://philpapers.org/rec/STEEIP-2>.
- Sterelny, Kim. 2012. *The Evolved Apprentice*. MIT Press.
- . 2021. *The Pleistocene Social Contract: Culture and Cooperation in Human Evolution*. Oxford, New York: Oxford University Press.
- Sweeny, Kate, Darya Melnyk, Wendi Miller, and James A. Shepperd. 2010. "Information Avoidance: Who, What, When, and Why." *Review of General Psychology* 14 (4): 340–53. <https://doi.org/10.1037/a0021288>.
- Tappin, Ben M., Gordon Pennycook, and David G. Rand. 2021. "Rethinking the Link between Cognitive Sophistication and Politically Motivated Reasoning." *Journal of Experimental Psychology: General* 150 (6): 1095. <https://doi.org/10.1037/xge0000974>.
- Thompson, Dennis F. 1993. "Understanding Financial Conflicts of Interest." *New England Journal of Medicine* 329: 573–573.

- Tollefsen, Deborah Perron. 2017. "Epistemic Reactive Attitudes." *American Philosophical Quarterly* 353–66.
- Tomasello, Michael, Alicia P. Melis, Claudio Tennie, Emily Wyman, and Esther Herrmann. 2012. "Two Key Steps in the Evolution of Human Cooperation: The Interdependence Hypothesis." *Current Anthropology* 53 (6): 673–92. <https://doi.org/10.1086/668207>.
- Van Bavel, Jay J., and Andrea Pereira. 2018. "The Partisan Brain: An Identity-Based Model of Political Belief." *Trends in Cognitive Sciences* 22 (3): 213–24. <https://doi.org/10.1016/j.tics.2018.01.004>.
- Wiersma, Miriam, Ian Kerridge, and Wendy Lipworth. 2018. "Dangers of Neglecting Non-Financial Conflicts of Interest in Health and Medicine." *Journal of Medical Ethics* 44 (5): 319–22. <https://doi.org/10.1136/medethics-2017-104530>.
- Wiersma, Miriam, Wendy Lipworth, Paul Komesaroff, and Ian Kerridge. 2020. "Symposium Lead Essay—Conflict of Interest: Opening Up New Territories." *Journal of Bioethical Inquiry* 17 (2): 169–72. <https://doi.org/10.1007/s11673-020-09987-6>.
- Williams, Daniel. 2021a. "Socially Adaptive Belief." *Mind & Language* 36 (3): 333–54. <https://doi.org/10.1111/mila.12294>.
- . 2021b. "Motivated Ignorance, Rationality, and Democratic Politics." *Synthese* 198 (8): 7807–27. <https://doi.org/10.1007/s11229-020-02549-8>.
- . 2021c. "Epistemic Irrationality in the Bayesian Brain." *The British Journal for the Philosophy of Science* 72 (4): 913–38. <https://doi.org/10.1093/bjps/axz044>.
- Zimmermann, Florian. 2020. "The Dynamics of Motivated Beliefs." *American Economic Review* 110 (2): 337–61. <https://doi.org/10.1257/aer.20180728>.