

[This is a pre-print of a paper appearing in *Res Philosophica*. Please cite the published version.]

## Epistemic Paternalism, Open Group Inquiry, and Religious Knowledge

Kirk Lougheed (University of Pretoria)  
philosophy@kirklougheed.com

### Abstract:

Epistemic paternalism occurs when a decision is made for an agent which helps them arrive at the truth, though they didn't consent to that decision (and sometimes weren't even aware of it). Common defenses of epistemic paternalism claim that it can help promote positive veritistic results. In other words, epistemic paternalism is often good for inquiry. I argue that there is often a better alternative available to epistemic paternalism in the form of what I call Open Group Inquiry. I then examine how Open Group Inquiry can be applied to cases of religious inquiry, while noting that epistemic paternalism is impermissible in cases of general religious inquiry. I argue that in the case of religious inquiry, there are serious questions about what constitutes evidence along with how to evaluate it. Rather than posing a particular worry for Open Group Inquiry, I suggest they pose a problem for religious inquiry in general. I conclude that while it very much matters how concepts like religious knowledge, religious faith, scepticism, etc., are defined that these considerations may well pave the way for a novel argument for religious scepticism.

### 1. Introduction

In the sciences and other related disciplines which prize positive veritistic outcomes some have argued that epistemic paternalism is sometimes permissible (Goldman 1991; Kristoffer Ahlstrom-Vij 2013). Epistemic paternalism can be defined, roughly, as: If agent X is going to make a doxastic

decision concerning question Q, and agent Y has control over the evidence that is provided to X, then, there are instances when Y need not make available to X all of the evidence relevant to Q if doing so will make X more likely to believe the truth about Q. Alvin Goldman (1991) argues that reflecting on certain real-life examples demonstrates that epistemic paternalism is sometimes (epistemically) permissible. The main reason for this is because there are cases where paternalism helps promote positive veritistic outcomes. For example, epistemic paternalism with respect to advertising, curriculum design in education, and rules of exclusion with respect to evidence in court is all motivated in an attempt to achieve positive veritistic results. Kristoffer Ahlstrom-Vij (2013) develops this line of defense even further by highlighting numerous cognitive biases that individuals are often unable to overcome on their own. Sometimes practising epistemic paternalism will allow individuals to avoid pernicious effects of various biases.

I have two main aims in this article. First, I examine whether epistemic paternalism is good for inquiry. Minimally, for a practice to be good for inquiry, it has to promote positive veritistic outcomes (i.e. help promote successful inquiry).<sup>1</sup> Initially one might think that given the veritistic defenses of epistemic paternalism it surely promotes successful inquiry. However, I argue that contra Goldman and Ahlstrom-Vij, epistemic paternalism is often impermissible in the very kinds of cases they describe because there is a better alternative. Thus, while I don't suggest that epistemic paternalism necessarily hurts inquiry in these cases, I believe there is a superior alternative practice. This alternative is called Open Group Inquiry. This occurs when individuals first inquire on their own before coming together in open dialogue about the correct epistemic evaluation. To support this model, I draw on arguments and information from the cognitive

---

<sup>1</sup> I'm going to use veritism throughout this paper, but I think what I say is consistent with promoting other positive outcomes such as justified belief, knowledge, understanding, etc.

division of labour, economic models about problem-solving and prediction, peer disagreement, and the interactionist account of reason.

My second aim is to explore whether my defense of Open Group Inquiry can also apply to cases of religious inquiry. I argue that it is indeed permissible and doesn't suffer from some of the problems uniquely associated with epistemic paternalism in the religious domain. However, I also argue that in the case of religious inquiry, there are serious questions about what constitutes evidence along with how to evaluate it. Rather than posing a particular worry for Open Group Inquiry, I suggest these questions pose a problem for religious inquiry in general. I conclude that while it very much matters how concepts like religious knowledge, religious faith, scepticism, etc., are defined that these considerations may well pave the way for a novel argument for religious scepticism.

## 2. Epistemic Paternalism Defended: It Promotes Positive Veritistic Results

Alvin Goldman (1991) argues that reflecting on a series of real-life examples demonstrates that epistemic paternalism is sometimes permissible. The main reason for this is because there are cases where paternalism helps promote positive veritistic outcomes. For example, epistemic paternalism with respect to advertising, curriculum design in education, and rules of exclusion with respect to evidence in court are all motivated in an attempt to achieve positive veritistic results. Goldman's discussion of the American legal system is particularly enlightening. He explains that the rules of evidence in the legal system are veritistic; they are aimed at creating the best possible epistemic

environment for the jurors to reach the correct verdict.<sup>2</sup> He writes that “the rules of evidence frequently foster the exclusion of evidence from the jurors [...] the fact that many of these rules are in place to avoid the jury being misled or unfairly prejudiced seems to amount to *epistemic* reasons” (Goldman 1991, 115-116).<sup>3</sup> We might interpret Goldman’s argument as endorsing the following conditional:

*If epistemic paternalism leads to the truth, then it is (epistemically) permissible.*

But Goldman needs something stronger in order to be justified since competing methods could be better at arriving at the truth. Consider:

*If epistemic paternalism leads to the truth more reliably than the relevant alternatives, then it is (epistemically) permissible.*

Let’s call this, the ‘Better than Alternatives Principle’ (or BAP). Goldman is careful to note he isn’t claiming BAP is universal in scope. For if it were universal in scope it would justify epistemic paternalism *every* time BAP obtains. Rather, Goldman says that “[e]pistemic paternalism on the part of isolated individuals is quite a different matter from paternalism exercised by the state, or any other powerful organ of society. There are historical reasons for being very cautious about state control of information” (Goldman 1991, 127). Likewise:

Epistemic outcomes obviously depend heavily on the (epistemically) good judgment of the message controller. If the controller accurately distinguishes true and false

---

<sup>2</sup> Goldman is writing from an American context, but I think his point applies to any adversarial legal system in liberal democratic societies.

<sup>3</sup> The motivation behind declaring certain pieces of relevant evidence inadmissible can be understood in terms of Bayesian probability. The worry is that the jury will assign the excessively high (or low) weight to the evidence and hence the relative probability of the defendant’s guilt (or innocence) will be updated (extremely) inaccurately (Goldman 1991, 117). In sum, “[t]he general idea is that the indicated rules of evidence are designed to protect jurors from their own ‘folly’, just as parents might keep dangerous toys or other articles away from children, or might not expose them to certain facts” (Goldman 1991, 118). He also offers additional justification for epistemic paternalism with examples from education curriculum decisions and rules surrounding advertising. In both cases, the paternalism has veritistic ends.

claims, and/or true and false pieces of evidence for claims (and their relevance), then the choice of included and excluded messages may well promote truth acquisition on the part of the audience (Goldman 1991, 124).

Thus, it's clear the scope of BAP isn't universal. However, Goldman doesn't come close to providing criteria informing us when epistemic paternalism is and isn't permissible. He doesn't tell us when BAP obtains or doesn't obtain (he also doesn't tell us when epistemic paternalism might be impermissible even in cases where BAP does obtain). He concludes that "when the cited variables or parameters take on appropriate values, epistemic paternalism will be justified. My previous examples are, quite plausibly, instances of this sort" (Goldman 1991, 125). We'll see below that Ahlstrom-Vij's more detailed defense of epistemic paternalism does provide criteria for when BAP obtains.

In his book, *Epistemic Paternalism: A Defence*, Ahlstrom-Vij builds on Goldman's veritistic defense of paternalism with an extensive discussion of biases and heuristics (2013). One way to understand Ahlstrom-Vij is that he is providing more details about when BAP obtains. He explains that heuristical reasoning is "roughly, reasoning that proceeds by way of sub-personal rules of thumb that operate on a limited number of cues, rather than through the systematic application of formal rules or principles of logic, statistics and probability theory [... and it] can lead to perfectly accurate results" (Ahlstrom-Vij 2013, 7).

Part of the reason why this empirical research supports epistemic paternalism is because we're unable to correct a number of problematic heuristical reasoning on our own. The first reason for this is because of the problem of motivation. This occurs because of the "bias blind spot which happens because we tend to underestimate the extent to which we, as opposed to everyone else, actually are prone to bias" (Ahlstrom-Vij 2013, 16). The second problem is about proper

correction. Even if we didn't suffer from the motivation problem "there are significant obstacles to not merely being motivated to correct for bias, but doing so successfully" (Ahlstrom-Vij 2013, 16). Ahlstrom-Vij concludes by suggesting that external constraints on reasoning are the best way to solve the problems he describes. The type of constraints Ahlstrom-Vij has in mind are, of course, paternalistic.<sup>4</sup>

Goldman never tells us what the criteria are for said variables and parameters (i.e. BAP obtaining), though he does tell us that examples like the evidential rules in court come close to meeting the invisible criteria he mentions. On the other hand, Ahlstrom-Vij helpfully provides jointly sufficient (though not necessary) conditions for when epistemic paternalism is permissible.

The two conditions are:

The alignment condition: The epistemic and non-epistemic reasons we have for implementing the paternalistic practice are not at odds with each other (Ahlstrom-Vij 2013, 117).

The burden-of-proof condition: The evidence indicates that it is highly likely that everyone affected by the paternalistic practice will benefit epistemically from it (Ahlstrom-Vij 2013, 122).<sup>5</sup>

The alignment condition is interesting, but it will not be my focus.<sup>6</sup> It's Ahlstrom-Vij's burden-of-proof condition that concerns me, and for the sake of argument I'm going to proceed as if it's all that's required for paternalism to be permissible. While the spirit of this condition is right

---

<sup>4</sup> A third defense of epistemic paternalism is distinct from the first two in not appealing directly to veritistic results. Pritchard suggests that another argument that could be employed in defense of epistemic paternalism is to suggest that epistemic goods are good simpliciter. The easiest route to this is to demonstrate that epistemic goods reduce to moral goods since moral goods are the most likely candidates to be goods simpliciter. One epistemic good that likely represents a moral good is understanding because it represents a strong cognitive achievement and hence represents virtue. I won't say more about this defense as it won't be my focus, though it's worth noting that Pritchard himself seems to think what he's about the matter is incomplete (Pritchard 2013).

<sup>5</sup> See Bishop 2014 for discussion.

<sup>6</sup> See Bullock 2018 for a critique of the alignment condition.

(and I think Goldman implicitly endorses something like it), I'm going to argue that this condition is (i) in fact is false and needs to be amended and that (ii) the amended condition shows why epistemic paternalism is rarely, if ever, permissible in the context of inquiry. For now, I'll briefly explain (i) and in the next section I'll defend (ii). The burden-of-proof condition really needs to be something like:

\*The burden-of-proof condition: The evidence indicates that it is highly likely that everyone affected by the paternalistic practice will benefit epistemically from it *and* there is no competing practice X which is incompatible with epistemic paternalism, where the evidence indicates that X is likely to yield better epistemic results for everyone than the paternalistic practice in question.

It is perhaps implicit in Ahlstrom-Vij's account that the burden-of-proof condition is meant to be consistent with \*the burden-of-proof condition. But notice what this amended condition makes explicit: if there's an epistemically better practice than paternalism, then paternalism isn't epistemically permissible. This is consistent with epistemic paternalism being better than *no* practice being enacted. This second formulation of the condition is compatible with BAP. Indeed, the problem with the first formulation is that it doesn't account for alternatives which is precisely BAP's focus. Finally, I take it that Ahlstrom-Vij wouldn't reject \*the burden-of-proof since it would be epistemically irrational to prefer one practice over another if it leads to worse epistemic results than the alternative.

### 3. What Is Inquiry?

Before presenting my alternative to epistemic paternalism it's important to first be clear about the precise meaning of inquiry itself. Jane Friedman explains that, minimally, inquiring is a certain kind of doxastic attitude. She says:

[T]rue inquirers are aiming to figure something out, but those others need not be. [For example, t]he detective is trying to figure out who robbed the bank and the reporter whether the working conditions are good. Their actions seem to count as part of their inquiries exactly because they are done in the service of their cognitive or epistemic aims or goals. A true inquirer then is someone with a certain kind of goal or aim, and so at the bottom of any true inquiry is a certain kind of aim- or goal-directed state of mind or attitude. This should come as no surprise – inquiry seems clearly to be an aim- or goal-directed activity (Friedman 2019, 298).<sup>7</sup>

Friedman continues to offer a much more robust account of inquiry and the associated doxastic states, but these aren't important for my purposes here. I will, however, proceed with the above definition of inquiry in view. It should be relatively uncontroversial and thus consistent with a wide variety of different understandings of the nature of inquiry.

One last distinction about inquiry is necessary before I continue. Friedman is clear that she thinks both individuals and groups can properly be thought of as inquirers (2019, 297). Strictly speaking Friedman is right that a single individual can conduct inquiry on her own. For that individual may be the only person who exemplifies the correct doxastic attitude toward some topic or question to rightly be considered an inquirer about it. But I'm going to focus on situations with two or more inquirers. I'm thus going to focus on group inquiry. This is, after all, the type of

---

<sup>7</sup> See also Friedman 2013 and 2017.

inquiry that most often occurs in the real-world. While it's possible to imagine a single isolated researcher who has little to no contact with the outside world, I'm not concerned about the best strategies for this sort of lone wolf inquirer.<sup>8</sup> Thus, in what follows when I refer to inquiry or research contexts, I refer to multiple inquirers. However, this distinction will admittedly be somewhat blurred in that I hold individuals must inquire in certain ways in order for group inquiry to ultimately succeed.

#### 4. Open Group Inquiry

In this section I argue that the antecedent of BAP doesn't obtain in the context of inquiry, because it doesn't lead to the truth more reliably than at least one relevant alternative. What I call Open Group Inquiry, a specific type of group inquiry, will usually outperform epistemic paternalism.

Consider that in order for an agent to reasonably believe the antecedent of BAP she needs to have good evidence that the paternalistic practice she wishes to engage in will yield positive veritistic results.<sup>9</sup> And in order for her to know that she needs to either (i) be aware of the truth of the matter in question or (ii) be aware of how some paternalistic practice is more likely to yield positive veritistic outcomes than other practices. This latter condition can be understood as suggesting that the agent needs to possess evidence that paternalism is the best method available, at least for her particular situation.

---

<sup>8</sup> Additionally, the recent focus on the social aspects of knowledge acquirement and transmission casts doubt on the possibility of a genuinely lone wolf inquirer in our world (even if one is logically possible).

<sup>9</sup> Of course, the antecedent of BAP could obtain even if an agent is not aware that it does so. But this is why there is a condition about reasonable belief. Externalists of a certain stripe may be sympathetic to the idea that in such a case it's perfectly legitimate rely on BAP even when agents aren't aware whether the antecedent obtains. I think this is misguided but won't disguise it further here. While I am trying to remain neutral between externalist and internalist accounts of justification the reasonable belief required pushes my argument, I suppose, explicitly towards internalism.

Here's the core of what I want to say about the permissibility of epistemic paternalism: the antecedent of BAP doesn't obtain in most cases of inquiry. Perhaps more accurately (and more modestly), I want to argue that cases of inquiry are most often situations where the balance of reasons suggest that the antecedent of BAP doesn't obtain. This is because there's a relevant and competing alternative that is usually better at leading to successful inquiry. This alternative is more likely than paternalism to lead to positive veritistic results. I call this alternative *Open Group Inquiry* and it is characterized by the following requirements and methodology:

*Open Group Inquiry*

- The group of inquirers is cognitively diverse. [Cognitive Diversity Requirement]
- All group members possess some degree of competence with respect to the specific inquiry at hand. [Competence Requirement]
- The individuals within the group of inquirers are given sufficient time and resources to examine the question *on their own*. [Individual Requirement]
- After examining the question on their own the inquirers come back together and discuss their solutions. [Group Requirement]
- The inquirers must be willing to share their results and listen to the others (without easily being moved from their position, but also without being impossible to convince otherwise). [Dialogue Requirement]
- The group is able to (eventually) reach a consensus. [Consensus Requirement]

These requirements are intended as necessary and sufficient conditions for Open Group Inquiry. They're the statement of what I take to be an epistemically superior practice to epistemic paternalism. In cases where these requirements can't be met (e.g. the group of inquirers happens to be cognitively homogenous and that cannot change), then epistemic paternalism might be

permissible. But when these requirements are (or can be) met, then Open Group Inquiry always trumps Epistemic Paternalism. A few additional details are in order.

One might wonder whether Open Group Inquiry is more successful when certain evidence is screened off from the group than when it is not screened off. However, this fundamentally misunderstands the nature of Open Group Inquiry. Part of inquiry itself is about determining *which* pieces of evidence are relevant and screening off evidence would not allow this process to play out. Another concern is that even if Open Group Inquiry does better with respect to multiple inquiries over time, it may not outperform cases of individual single shot inquiries where epistemic paternalism is employed. In other words, there's a difference between comparing multiple inquiries over time versus just one inquiry. As I will attempt to show below, my claim is that there are forward-looking reasons to *almost always* prefer Open Group Inquiry with respect to multiple inquiries and just single inquiries (at least when Open Group Inquiry is possible). Epistemic paternalism could, in specific cases, perhaps outperform Open Group Inquiry. While this is possible, it's unlikely and since there is little reason to think this will happen there is little reason to choose epistemic paternalism over Open Group Inquiry when deciding which method of inquiry to pursue.

Finally, I'm not making any claims about *how* to enact Open Group Inquiry. Perhaps it is exceedingly difficult to ensure that inquirers aren't too conciliatory or too stubborn. Maybe there is not enough time or resources to conduct the appropriate inquiry. I am only trying to lay the theoretical framework for Open Group Inquiry, rather than address these practical considerations. I am merely claiming that *if* it is enacted it will almost always outperform epistemic paternalism. This means that the antecedent of BAP rarely, if ever, obtains. There is, I think, ample

philosophical and empirical research to support my position. In what follows I briefly examine some of the supporting evidence for Open Group Inquiry.

#### 4.1 Problem Solving And Prediction

The first evidence for Open Group Inquiry I examine is based on the idea that groups made up of heterogeneous individuals are always better at problem-solving and prediction than homogenous groups made up of more individually talented people. Economist Scott Page argues persuasively that in many scenarios, diverse groups are better at problem solving and prediction making than homogenous groups. This means that individuals are best left alone to pursue their own lines of inquiry, free from paternalistic practices. Page says that he “will show with modest rigor how diverse perspectives, heuristics, interpretations, and mental models improve our collective ability to solve problems and make accurate predictions” (Page 2004, xx). Individual members of homogenous groups are going to understand problems in the same way. Inasmuch as one member finds a problem intractable, every member of the group will find it so. Groups with members who approach problems from a wide variety of perspectives are more likely to discover solutions to problems that initially appear intractable (Page 2004, 52-53). According to Page, “collections of diverse individuals outperform collections of more individually capable individuals” (Page 2004, 159) when the following four conditions are met: (1) the problem needs to be difficult such that no single group member always discover the best or preferred solution to the problem; (2) the reasoners need to be intelligent and able to communicate with each other effectively; (3) there needs to be sufficiently diverse solutions proposed and; (4) there needs to be a large pool of problem solvers to draw upon (Page 2004, 162). While these four conditions aren't the only ones

that when met will entail that diversity trumps ability, they do guarantee it. One might protest that none of these ideas conflict with epistemic paternalism. Yet consider that one of the functions of paternalism is to have every member of the group think along the same lines. This is clear in many of the examples Goldman offers such as educational policy, advertising, and the evidential rules of exclusion in court cases.<sup>10</sup>

#### 4.2 The Epistemic Benefits Of Peer Disagreement

The epistemology of disagreement asks whether an agent is rational to continue to believe P when she discovers an epistemic peer – a cognitive and evidential equal – who believes not-P. Some have argued that it's rational for her to continue to belief P in the face of disagreement because there are epistemic benefits in the offing if one remains steadfast (Elgin 2010, 2018; Kraay forthcoming; Lougheed 2020). Consider the cases of Semmelweis and handwashing and Darwin and natural selection. They each faced considerable disagreement at the time they made their respective claims yet remained steadfast. And we're epistemically better off because they did so (Lougheed 2020, xx). In relevantly similar situations, then, epistemic peers are justified to remain steadfast in the face of disagreement. If epistemic peers conciliated in the face of disagreement, then groups of inquirers would lose cognitive diversity which is essential for successful inquiry.<sup>11</sup>

#### 4.3 The Interactionist Account Of Reason

---

<sup>10</sup> Page's work is the subject of recent controversy, though much of the debate seems to be regarding the scope of his claims. Likewise, if Page is wrong the evidence for my preferred account is weakened, but not defeated. See Thomson 2019 for criticism of Page. See Singer forthcoming for a reply.

<sup>11</sup> There are many details about this type of argument which I don't have the space to discuss here. See my Lougheed 2020 for a book length defense of this argument.

Hugo Mercier and Dan Sperber are well-known for defending what they call the ‘Interactive’ approach to reasoning (2017).<sup>12</sup> This runs directly contrary to the dual-systems approach which has it that humans use two systems when reasoning, type 1 (fast) and type 2 (slow).<sup>13</sup> Mercier and Sperber claim their theory can better explain a whole host of data than the dual-systems account. Consider that we are bad at producing reasons. When we do, we often exhibit a whole host of biases. We exhibit a kind of lazy ‘myside’ bias. However, it turns out that we are actually quite good at evaluating the arguments of others, even if we don’t evaluate our own very well. This isn’t surprising. We don’t need to waste cognitive resources developing arguments since they’re going to be scrutinized by others later. Additionally, it’s an adaptive advantage to only trust the right people (i.e. those who trustworthy). Thus, it makes sense we’re adept at evaluating the reasons of others. This interactive approach to reason has it that we are at our bests when we are arguing and reasoning *with others*. The idea that we reason well on our own provided we have enough time is, according to Mercier and Sperber, simply a myth. We reason better when we reason with others. But this doesn’t reject the Individual Requirement. Instead, it shows why the Group and Dialogue Requirements need to be met after it.

## 6. Epistemic Paternalism And Religious Knowledge

Below I’m going to argue that Open Group Inquiry can successfully be applied to cases of religious inquiry before concluding with some general questions about religious inquiry. However, I’m going to briefly explain why epistemic paternalism is often impermissible in cases of religious

---

<sup>12</sup> See also Mercier and Sperber 2011a, 2011b.

<sup>13</sup> See Kahneman 2013.

inquiry even if Open Group Inquiry weren't an alternative. To see why this is the case consider the following example.

Suppose that Sally is curious about religion. Further suppose that she was raised in a devout religious home but has many doubts and questions. Specifically, she wants to inquire into whether monotheism is true, or at least epistemically rational to believe. Here's why epistemic paternalism couldn't be used to help Sally discover religious truths: In order for her to benefit from another agent acting paternalistically towards her with respect to religious knowledge she would first need to know which religious authority to trust. And there are numerous religious authorities who issue competing verdicts. But since Sally wants to discover whether monotheism in general is true, part of the very question she is asking is which religious authority she can trust. It's true that in very specific religious inquiries paternalistic practices could perhaps be helpful. If Sally already happened to be a committed Catholic who simply wanted to learn more about the Trinity, there are clear authorities she could trust on such matters (though it's not as if there isn't intra disagreement on many such topics too). Thus, while there might be instances where epistemic paternalism is helpful regarding very specific religious questions for those who are already settled on much of what they believe, it is of little help in a general religious inquiry like Sally's.<sup>14</sup>

## 7. Open Group Inquiry And Religious Knowledge

---

<sup>14</sup> One might object that *if* Sally's community is right (say, for example, that she was raised in a Christian community), then epistemic paternalistic practices employed by that community would indeed serve her better than Open Group Inquiry. Again, this is likely a place where I am committed to internalism. In order to know this about her community then Sally would need reasons prior to her investigation for thinking that she should trust her community. But whether she should trust them is part of what's at stake in the first place.

There is no *prima facie* reason for thinking Open Group Inquiry cannot be successfully applied to cases of religious inquiry. To begin, religious inquiry is compatible with my general description of inquiry taken from Friedman. A particular individual or group can spend a specific amount of time inquiring into questions such as whether God exists, the Trinity is coherent, or Jesus rose from the dead, etc. We might think inquiry of this sort is quite different from scientific inquiry and indeed it is (though empirical evidence is not ruled out for certain religious questions). Nonetheless, this isn't a reason to exclude pursuing answers to various religious questions as a genuine form of inquiry. To do so would be to exclude a vast array of non-empirical fields as legitimate fields for inquiry. Such a position is too extreme and sounds suspiciously close to endorsing some version of logical positivism or an unpalatable scientism.<sup>15</sup> An easy way to explain how Open Group Inquiry could apply to religion is by way of example.

Recall our friend Sally who wants to investigate the truth of monotheism. Sometimes, it seems that this sort of religious inquiry is often pursued by individuals alone. Sally needs to discover religious truths on her own and for herself. However, inasmuch as this is even possible (which is doubtful), Open Group Inquiry gives us reason to doubt that 'going it alone' is the most effective approach for Sally to take. Rather, Sally should use Open Group Inquiry to seek out religious truth. This means she needs to conduct her inquiry alongside of a group of cognitively diverse inquirers. Presumably, this means that there needs to be a certain degree of religious diversity in the group. A group of religious inquirers who all ascribed to the same religion are more likely to experience groupthink, get stuck in the same places, and fail to see novel solutions, etc. Furthermore, each member of Sally's group needs to have a certain degree of expertise with respect

---

<sup>15</sup> See Ayer 1971.

to religion. I won't stipulate the level required here, but presumably the group would be better off if each member had graduate level training in a wide variety of relevant subjects.

Sally and her group members also need sufficient time and resources to conduct religious inquiry on their own before coming together as a group.<sup>16</sup> Again, I'm not going to stipulate an answer here, however, Sally needs to be open to the idea that this process could take a very long time. When each individual in the group member has concluded their inquiry, they must be willing to come together and share their results in a productive manner. This means being able to effectively share their results while listening to other members. The group needs to dialogue and to eventually reach a consensus.

I don't claim that this is a perfect example, but it at least shows that Open Group Inquiry can be plausibly applied to religious inquiry. There are, of course, a number of objections, some of which I will address in the next section.

## 8. Objections To Open Group Inquiry In Religious Contexts

### 8.1 Can Epistemic Paternalism Promote Open Group Inquiry?

One objection to Open Group Inquiry asks whether paternalism could be used to promote cognitive diversity. Suppose that you and I both believe *P*, but I know that cognitive diversity is epistemically beneficial.<sup>17</sup> So further suppose that I employ a paternalistic practice in order to get you to believe *not-P*. This doesn't promote consensus in the short-term, but it might help us to reach a consensus

---

<sup>16</sup> That we are very earlier in our religious inquiry as a species has been recently emphasized in a number of books and articles by J.L. Schellenberg (2013, 2019). He explains that religious inquiry is extremely young relative to the first homo sapiens. Likewise, we are set to live another billion years (pending climate and/or nuclear disaster). Schellenberg says that in light of the deep future there may be evidence and indeed solutions to problems we can't even fathom. Page clearly doesn't have the idea of the deep future in view when he says problem solvers need to be given sufficient time (he seems to mean some finite amount of time within a normal human lifespan). However, if Schellenberg's account is right this emphasis on sufficient time is even more appropriate.

<sup>17</sup> See Fleisher forthcoming.

about the truth of  $P$  in the long-term. Since this scenario is possible, the worry is that epistemic paternalism is at least as good as Open Group Inquiry at promoting positive veritistic results.

*Reply:*

There are at least three responses to this worry. The first is that this wouldn't show that epistemic paternalism is *better* than Open Group Inquiry. Indeed, that epistemic paternalism could be used to promote some of the requirements of Open Group Inquiry is perfectly consistent with epistemic paternalism being epistemically worse than Open Group Inquiry *on its own*. Second, I haven't argued that epistemic paternalism is *always* impermissible. In cases of inquiry where Open Group Inquiry is impossible then epistemic paternalism may well be the best way to promote positive veritistic outcomes. That epistemic paternalism might, in certain cases, be able to promote Open Group Inquiry could simply turn out to be another instance where it's permissible. Again, this is perfectly consistent with everything I've argued for thus far. Third, that epistemic paternalism could be used to promote cognitive diversity doesn't entail that it could be used to help all of the other requirements of Open Group Inquiry obtain. It thus doesn't follow that because it could be used to help the Cognitive Diversity Requirement obtain that it can be used to help Open Group Inquiry in general obtain. A separate argument for that conclusion would have to be provided by the proponent of this worry.

## 8.2 Are There Better Practices?

Even granting that Open Group Inquiry is superior to epistemic paternalism with respect to arriving at religious knowledge, are there better epistemic practices with respect to acquiring religious knowledge? Perhaps there are epistemic practices unique to acquiring religious knowledge (e.g. religious experience or religious practice).

*Reply:*

A full answer to this question involves knowing how to differentiate between various epistemic practices in addition to knowing how to evaluate those practices. Given the large nature that such a project would take I can't pursue this type of reply here. However, it's important to note that one of the claims of Open Group Inquiry is that it will necessarily outperform homogenous groups or lone wolf inquirers. This in itself eliminates a large number of other methods of inquiry from consideration. In light of this I suggest that it's likely Open Group Inquiry is one of the very best methods by which to inquire about religion.

### 8.3 Religious Knowledge Is A Uniquely Individualistic Enterprise

A perhaps especially protestant objection to Open Group Religious Inquiry is that religious knowledge ought to be arrived at alone. This isn't to say that one can't have help from others in arriving at such knowledge. Rather, it's wrong to *depend* on others for religious knowledge. Someone like Sally ultimately needs to come to religious knowledge of her own accord.

*Reply:*

I take it that such an objection demonstrates a naivety about the nature of knowledge acquisition in the first place. Over the past few decades, epistemologists have increasingly recognized the social dimensions of knowledge. Topics that have garnered considerable attention include testimony, disagreement, and epistemic injustice. Other topics like the nature of group knowledge and epistemic paternalism are beginning to receive more attention too. To think we ever acquire knowledge in a vacuum is not only naïve, but false.

### 8.4 Only God Can Provide Genuine Religious Knowledge.

Another way to motivate an objection similar to the one above is based on the idea that religious inquiry can't be successful without the cooperation of God. In some Christian traditions it is held that a full understanding of the gospel is only possible via the testimony of the Holy Spirit (the

third person of the trinity). In the case of coming to religious knowledge the epistemic practice that is used is not really important. For no matter which doxastic practice is employed it can't be successful without the appropriate divine intervention. Notice that this objection nicely sidesteps worries about gaining religious knowledge in a vacuum.

*Reply:*

This type of objection raises a number of interesting questions. One such question is whether religious knowledge is equivalent to religious faith. In loose everyday talk these two concepts are often used interchangeably, but it's far from clear that they're identical. Consider that it's consistent to say that someone knows quite a lot about religion but doesn't have any religious faith. Maybe the worry is that an agent could come to religious knowledge by employing Open Group Inquiry but not genuine religious faith. Additionally, discussions about the nature of faith are on the rise in analytic philosophy of religion. Distinctions have been made between faith *in* God versus having faith *that* God. Furthermore, non-doxastic accounts of faith are growing in popularity: Perhaps we can accept, plan, suppose, or act as if God exists even if it's epistemically irrational to *believe* that God exists (e.g. Alston 1996; Howard-Snyder 2017; Pojman 1986). It seems that at least in some cases, part of the motivation behind such non-doxastic accounts of faith is due to scepticism about our ability to have religious knowledge. If scepticism towards religious knowledge is correct, then it doesn't matter what epistemic practice is in use, humans can't (or are unlikely to) possess religious knowledge. On this view Open Group Inquiry is incompatible with non-doxastic accounts of faith that are motivated by scepticism. This is an interesting result in itself, but especially in light of the growing popularity of such accounts.

8.5 Religious knowledge acquired this way cannot result in religious faith

Another similar objection is in reference to Pascal's wager. Indeed, Pascal was aware that we can't will ourselves to believe that God exists. So, he recommended that we should just continually put ourselves in positions to form such beliefs. However, some have objected that arriving at faith in this way is disingenuous. Thus, even if Open Group Inquiry turned out to be compatible with non-doxastic accounts of faith, questions remain.

*Reply:*

All that this objection would demonstrate is that while it might be possible to arrive at religious knowledge via Open Group Inquiry, it's impossible to use it to arrive at religious faith. Again, this is an interesting result, but not one that hurts the plausibility of Open Group Inquiry.

8.6 Does Open Group Inquiry really recommend taking a poll to arrive at religious knowledge?

Some might worry that Open Group Inquiry seems to recommend the bizarre idea of simply taking a poll in order to arrive at religious knowledge. However, it's clear that taking a poll isn't a good way to arrive at religious knowledge. So Open Group Inquiry cannot be applied to cases of religious inquiry.

*Reply:*

A poll is not what is recommended by Open Group Inquiry. Not just anyone can participate in the religious inquiry conducted by Sally. Inquirers have to be sufficiently competent, able to express their findings, etc. Likewise, Open Group Inquiry does not recommend polling the group. The group is supposed to reach consensus through dialogue. This is a very different process from conducting a poll.

8.7 Can A Group Of Religious Inquirers Ever Reach Consensus?

A related worry which touches on something quite important asks whether a group could really ever reach consensus about religious matters. On some subject matters there are (relatively) clear

and generally agreed upon criteria for what constitutes evidence along with how to evaluate that evidence. This doesn't mean that there is complete universal agreement about evidence and evaluation amongst practitioners in the subjects in question. However, consider a mathematician who claims to have solved a proof that no one else to date has been able to solve. There is consensus amongst the mathematical community about what would constitute decisive evidence that the proof is sound. Sometimes it can take many months or years for complicated proofs to be verified by the mathematical community (though artificial intelligence is showing signs of speeding this process up). Likewise, in the sciences it is often quite clear what counts as confirming or disconfirming evidence. Such evidence is even sometimes obvious to laypeople. That a certain plane design will be able to fly or that a bridge design will be able to support cars are easy claims to check.

This worry need not take the form of assuming scientism or verification. It need not assume that anti-realism is true in non-empirical fields like religion. For there could well be objective truths about religious matters even if they can't be discovered by Open Group Inquiry. Rather, the question here is about whether it's right to expect Open Group Inquiry to succeed, especially regarding the consensus requirement given how little agreement there is about what constitutes evidence in addition to how to evaluate it in the religious domain. If one doubts this, consider the following question: what is the evidential import of testimony from a generally reliable person about a religious experience she had? Consider that opinions about how to treat testimony about religious experiences vary widely, even within religious communities themselves.

*Reply:*

I don't claim to have a definitive answer to this worry. However, I do have two observations: The first is that worries about what constitutes evidence and how to assess it in the case of religion

aren't just problems for Open Group Inquiry; they're problems for *any* method of religious inquiry.<sup>18</sup> Second, this objection might actually point towards religious scepticism in general rather than to a particular fault with Open Group Inquiry. In other words, the fact that Open Group Inquiry might not be able to apply to religious inquiry should actually give us pause about our ability to acquire religious knowledge (especially given that many of the objections here apply to all methods of religious inquiry too). Of course, some of this will very much depend upon how one defines religious knowledge, including whether it is taken to be conceptually distinct from religious faith. I leave such questions for another day.

#### 9. Averroes On Epistemic Paternalism And Open Group Inquiry

Before concluding, it is an interesting exercise to examine how certain philosophers might receive my defense of Open Group Inquiry. Elsewhere, Joshua Harris and I argue that the Islamic medieval philosopher Averroes clearly endorses what ought to be considered epistemic paternalism, particularly in religious settings (Lougheed and Harris Manuscript). It's thus instructive to explore how Averroes might react to Open Group Inquiry.

Harris and I explore a passage from the *fasl al-maqal* where Averroes appears to be defending epistemic paternalism. In the passage in question, an unnamed woman announces that God is 'in the sky' and this is approved of by The Prophet (*FM*, 45).<sup>19</sup> However, God is technically not *in* the sky or heaven for that matter; God has no spatio-temporal location as such. Yet the

---

<sup>18</sup> Again, this probably commits me to internalism. Reliabilist theories that are externalist do not require that the knower in question be in possession (or be able to access) the evidence for the belief in question.

<sup>19</sup> Citations of the *Fasl al-maqal* are drawn from Averroes, *Decisive Treatise*, trans. G.F. Hourani, reprinted in J. McGinnis and D.C. Reisman (eds.), *Classical Arabic Philosophy: an Anthology of Sources* (2007, 309–30). In-text citations use the abbreviation "*FM*," along with a section number.

Prophet credits the woman with having faith. Of course, Averroes isn't going to say that The Prophet or the *hadith* is incorrect, nor can he hold that a clearly false belief is a sign of faith.<sup>20</sup> Instead, he distinguishes between knowledge by imagination and knowledge by demonstration. For Averroes, knowledge by demonstration (*burhan*) is akin to scientific and/or literal understanding.<sup>21</sup> Knowledge by imagination (*tasdiq ala' qabil min tahayyul*) is *not* scientific and is instead non-literal and metaphorical (*FM*, 54).<sup>22</sup> In the case in question, then, the woman has knowledge via imagination instead of knowledge by demonstration. This is how she can be said to have genuine religious faith.

Averroes' defense of epistemic paternalism appears to apply particularly strongly in religious contexts since the stakes are so high. Consider that some people are (through no fault of their own) not in a position to understand a demonstration about why God does not literally have a spatio-temporal location. Attempting to explain this to the woman with a demonstration this could cause confusion and ultimately undermine her faith (Lougheed and Harris Manuscript). So, it is best to use imagination and preserve her faith. Harris and I argue that the distinction between demonstration and imagination is one unexplored way to justify epistemic paternalism in contemporary cases, especially when the stakes of the belief in question are high (Manuscript).<sup>23</sup> It's thus clear that in at least certain cases Averroes endorses something like contemporary epistemic paternalism.<sup>24</sup>

---

<sup>20</sup> For discussion see Lougheed and Harris Manuscript and Faizi 2016.

<sup>21</sup> In following Aristotle knowledge by demonstration for Averroes is about discovering causal powers. See Lougheed and Harris Manuscript.

<sup>22</sup> Harris and I further explain that knowledge via imagination tends to be more closely associated with spiritual knowledge. See Lougheed and Harris Manuscript and *EDA*, 62. *EDA* refers to Deborah Black's (2009) unpublished translation of Averroes' *Epitome of the De anima*, which can be found at <http://individual.utoronto.ca/dlblack/translations.html>.

<sup>23</sup> One upshot of our view, then, is that this justification extends beyond the religious domain.

<sup>24</sup> For a full defense of this claim see Lougheed and Harris Manuscript.

What might Averroes say about Open Group Inquiry in religious settings? I don't see any reason to think that Averroes would be *in principle* against Open Group Inquiry. If the woman in question was able to practice Open Group Inquiry, then she would be likely to arrive at the truth. Let's assume for a moment that Open Group Inquiry is aiming at knowledge by demonstration. Where Averroes might be sceptical is whether the woman in question is in a position to conduct Open Group Inquiry in the first place. Remember that Open Group Inquiry doesn't say that just *any* heterogenous group of inquirers will do. Rather, each member of the group in question must meet a minimum level of competence with respect to the issue in question. Averroes might wonder whether laypersons have the requisite level of competence to practice Open Group Inquiry about religious matters. If this turned out to be the case, then BAP would obtain and the sort of epistemic paternalism he recommends would be reasonable to employ.

A further line of inquiry that I can only briefly touch on here regards the assumption that when translated into contemporary terms, Open Group Inquiry is aimed at knowledge via demonstration. But is this assumption required? Could Open Group Inquiry be successfully practiced if the aim were instead knowledge via imagination? In attempting to answer such questions it's noteworthy that all (or almost all) contemporary defenses of epistemic paternalism are *veritistic*. They appeal to the fact that epistemic paternalism will, in certain cases, better help agents to arrive at the truth (see Section 2 above). In proposing my alternative to epistemic paternalism in Open Group Inquiry, I have claimed that it is superior (when certain conditions obtain) at arriving at the truth. In light of this fact, then, it's at least *prima facie* plausible that Open Group Inquiry could be employed to help arrive at the truth via imagination. The women in the

example that Averroes is analyzing is, after all, thought to possess religious truth (though it is non-literal).<sup>25</sup>

Finally, one might wonder whether Averroes is not really in principle opposed to Open Group Inquiry in religious settings. He might well say that since The Prophet possess religious truth then he always ought to use paternalistic practices to ensure that laypersons keep their faith. Yet the perspective from which we're evaluating this claim very much matters. From the perspective of the religious leader who is convinced that their religious beliefs are correct, they may well be justified in employing epistemic paternalistic practices to ensure the laypersons under their care have the appropriate religious faith. However, the point remains that this isn't epistemically best from the *perspective of the inquirer* who doesn't have a settled position on the subject matter in question and is attempting to figure out the best way engage in the inquiry.<sup>26</sup>

## 9. Conclusion

Veritistic defenses of epistemic paternalism are only successful when the antecedent of BAP obtains. I've argued that the antecedent rarely, if ever, obtains in research contexts because Open Group Inquiry is a superior alternative. Heterogenous groups outperform homogenous groups of more individually talented individuals at tasks like problem solving and prediction. Remaining steadfast in the face of peer disagreement maintains cognitive diversity and thus leads to epistemic benefits. The interactionist account of reason provides additional support for the idea that reasoning works best when individual inquirers remain in dialogue with one another. I further

---

<sup>25</sup> There are important questions here regarding the conception of truth required for this account to succeed but I do not have the space to address this issue here.

<sup>26</sup> Whether the justification for epistemic paternalism is subject-relative in this way is yet another area of inquiry that remains to be addressed in the literature.

argued that epistemic paternalism is often impermissible in cases of religious inquiry because which religious authority to trust is often part of the very inquiry itself. I then suggested that Open Group Inquiry can plausibly be applied to cases of religious inquiry. There are practical challenges to enacting it in real-life scenarios, but there is no in principle reason why it could not be employed. I concluded by addressing a number of objections to Open Group Religious Inquiry. The most pressing objections ask about the implications of the fact that there is little consensus about what constitutes evidence in the religious domain, in addition to how to assess such evidence. While this is a legitimate challenge to Open Group Religious Inquiry, it's really a challenge to religious inquiry in general. It may well point towards general religious scepticism, though much more would have to be said to defend that conclusion. Finally, I concluded by examining whether Averroes would approve of Open Group Inquiry since he clearly seems to endorse epistemic paternalism. I tentatively suggested that there's little reason to think Averroes would be in principle against it. However, it the subject in question very much matters along with the perspective from which the best strategy to inquiry is being asked.<sup>27</sup>

---

<sup>27</sup> Thanks to Billy Dunaway, Joshua Harris, and Jon McGinnis. Work on this paper was supported by funding from the John Templeton Foundation.

## Bibliography

- Ahlstrom-Vij, Kristoffer. (2013). *Epistemic Paternalism: A Defence*. London: Palgrave Macmillan.
- Alston, William P. (1996). "Belief, acceptance, and religious faith." In *Faith, Freedom, and Rationality*, edited by Jeff Jordan and Daniel Howard-Snyder, 3-27. Lanham MD: Rowman & Littlefield.
- Averroes. (2009). *Epitome of the De Anima*, trans. Deborah L. Black. Unpublished manuscript available at <http://individual.utoronto.ca/dlblack/translations.html>.
- Averroes. (2007). *The Decisive Treatise*. In *Classical Arabic Philosophy: An Anthology of Sources*, edited by Jon McGinnis and David C. Reisman, 309–330. Indianapolis: Hackett Publishing Company.
- Ayer, A.J. (1971). *Language, Truth and Logic*. New York: Penguin.
- Bishop, Michael A. (2014). Review of Kristoffer Ahlstrom-Vij (2013). *Epistemic Paternalism: A Defence*. Palgrave MacMillan. In *Notre Dame Philosophical Reviews*.
- Bullock, Emma. (2018). "Knowing and Not-Knowing For Your Own Good: The Limits of Epistemic Paternalism." *Journal of Applied Philosophy* 35 (2): 433-447.
- Elgin, Catherine. (2018). "Reasonable Disagreement." In *Voicing Dissent: The Ethics and Epistemology of Making Disagreement Public*, edited by Casey Rebecca Johnson, 10-21. New York: Routledge.
- \_\_\_\_\_. (2010). "Persistent Disagreement." In *Disagreement*, edited by Richard Feldman and Ted A. Warfield, 53-68. Oxford University Press.
- Faizi, N. (2016). "Averroes on the Relationship between Philosophy and Scripture: A Conditional Hierarchy." *The Journal of Scriptural Reasoning* 15 (1): 1-14.

- Fleisher, Will. (forthcoming). "Rational Endorsement." *Philosophical Studies*.
- Friedman, Jane. (2019). "Inquiry and Belief." *Nous* 53 (2): 296-315.
- \_\_\_\_\_. (2017). "Why Suspend Judging?" *Nous* 51 (2): 302-326.
- \_\_\_\_\_. (2013). "Suspended Judgment." *Philosophical Studies* 162: 165-181.
- Goldman, Alvin. (1991). "Epistemic Paternalism: Communication Control in Law and Society." *The Journal of Philosophy* 88 (3): 113-131.
- Howard-Snyder, Daniel. (2017). "Markan faith." *International Journal for Philosophy of Religion* 81: 31-60.
- Lougheed, Kirk. (2020). *The Epistemic Benefits of Disagreement*. Springer.
- Lougheed, Kirk and Joshua Harris. "Epistemic Paternalism, Averroes, and Religious Knowledge." Unpublished Manuscript.
- Mercier, Hugo and Dan Sperber. (2017). *The Enigma of Reason*. Massachusetts: Harvard University Press.
- Mercier, H. (2011a) "Reasoning Serves Argumentation in Children." *Cognitive Development* 26 (3): 177-191.
- \_\_\_\_\_. (2011b). "Why do humans reason? Arguments for an argumentative theory." *Behavioral and Brain Sciences* 34 (2): 57-74.
- Page, Scott. (2007). *The Difference: How the Power of Diversity Creates Better Groups, Firms, and Societies*. USA: Princeton University Press
- Pojman, Louis. (1986). "Faith without belief." *Faith and Philosophy* 3: 157-176.
- Pritchard, Duncan. (2013). "Epistemic Paternalism and Epistemic Value." *Philosophical Inquiries* 1 (2): 9-37.
- Kahneman, Daniel. (2013). *Thinking, Fast and Slow*. Anchor Canada.

Schellenberg, J.L. (2019). *Religion after Science: The Cultural Consequences of Religious Immaturity*. Cambridge University Press.

\_\_\_\_\_. (2013). *Evolutionary Religion*. Oxford University Press.

Singer, Daniel J. (2019). "Diversity, Not Randomness, Trumps Ability." *Philosophy of Science* 86 (1).

Thomson, Abigail. (2014). "Does Diversity Trump Ability? An Example of the Misuse of Mathematics in the Social Sciences." *Journal of the AMS* 61 (9): 1024-1030.