

Religious Belief and the Wisdom of Crowds

Jack Warman and Leandro De Brasi
University of La Frontera

The final version of this article is forthcoming in *Sophia*. Please cite the published version.

§1 Introduction.

In their simplest form, *consensus gentium* arguments for theism argue that theism is true on the basis that everyone believes that theism is true. While such arguments may have been popular in history, they have all but fallen from grace in the philosophy of religion. In this short paper, we reconsider the neglected topic of *consensus gentium* arguments, paying particular attention to the value of such arguments when deployed in the defence of theistic belief. We argue that while *consensus gentium* arguments are unlikely to offer anything close to overwhelming support for theism, their probative value is nevertheless underappreciated, and they have been unfairly maligned as a consequence.

§2 *Consensus gentium* arguments and their discontents.

Consensus gentium arguments receive surprisingly little attention in contemporary Western philosophy of religion. (We'll discuss some notable exceptions later.) As Kelly suggests, this is in part because philosophers are more interested in constructing and evaluating arguments, and less accustomed to citing opinion polls (2011). In the event that philosophical debate concerns itself with public opinion, it is often because the former has the latter in its sights, so to speak.

In their most basic form, *consensus gentium* arguments involve an inference from a fact about what people believe (usually a majority thereof) to a conclusion about what is in fact the case. In the history of philosophy, such arguments are made in defence of theism. Here's an extremely simple—and clearly flawed—*consensus gentium* argument:

- (Premise 1) Everyone believes that theism is true.
- (Conclusion) Therefore, it is the case that theism is true.

Premise 1 appears to be empirical: it makes a claim about the mental states of all the people in the world. The first premise is also, of course, false. It is not the case that everyone believes that theism is true. Atheists believe that theism is false, and agnostics withhold judgment regarding theism. There are added complications when we consider how polytheists, pantheists, and deists fit into this picture.

Now, one might attempt to fix the first premise. One option, endorsed by Joyce (1923), is to argue that while there are some atheists 'out there', the consensus is that God exists: most people are theists, and apparent instances of atheism can be explained away as isolated (and thus irrelevant) deviances from the norm. We can amend the first premise accordingly, which leaves us with the following version of the *consensus gentium* argument:

- (Premise 1*) Most people believe that theism is true.
- (Conclusion) Therefore, it is the case that theism is true.

In this version of the *consensus gentium* argument, premise 1 has a better chance of being true. A *prima facie* case can be made that most people in the world believe that theism is true. To be more precise,

since most people don't talk about these things in such technical language, it is probably more likely that most people hold beliefs which, if true, would imply that theism is true. We can imagine someone who wouldn't identify as a theist using that term but who nevertheless claims that 'there must be some higher being out there, watching over us, beyond the limits of our perception'. That should be sufficient for the purposes of one who wishes to defend this kind of *consensus gentium* argument.

Maybe total unanimity is more than just an unrealistic ideal: one might think that total unanimity about theism comes with problems of its own. Some recent work on the so-called *paradox of unanimity* can guide us here (Gunn *et al.* 2016). The paradox of unanimity arises from the apparent logical incompatibility of the following claims:

- A. *Unanimity is evidence of success.* When a deliberating group votes whether p , unanimity is the ideal: as a group's judgement converges on unanimity, the more confident we can be in the deliberating group's judgements.
- B. *Unanimity is evidence of error.* In a significant number of cases, unanimity is so unlikely that it should be taken as evidence of error.

Why does it seem that unanimity is evidence of success? Suppose you want to confirm whether Tokyo is the capital city of Japan before writing a paper in which you use that city in a thought experiment. You ask one of your colleagues. They tell you that it is. You ask another colleague. They agree. Still not content, you walk around the department, asking every colleague that you find whether Tokyo is indeed the capital of Japan. They all confirm what the first colleague told you. It is unanimous, and that seems like evidence of success, which is to say, evidence that Tokyo really is the capital of Japan. At first you only had a little bit of evidence—one person's testimony—but as you asked your colleagues you gained more and more evidence, until you ended up with abundant evidence that Tokyo is Japan's capital city. Unanimity seems to be evidence of success because the more evidence you have that p , the more confident you can be that p . (Of course, it's worth noting that unanimity isn't the only factor to consider – there's a reason you sought multiple opinions rather than making up your mind after asking one person. More on that later.)

Why does it seem that unanimity is evidence of error? Suppose you want to find out whether the authoritarian leader of a country is popular among the citizens. You decide to check the results of the previous election in that country, and to your surprise, they won 100% of the vote. Not a single person voted for any other candidate, and nobody spoiled their ballot. This is very surprising to you. But of course, the reason it's surprising is not because you hadn't realised that the leader was so popular. Rather, you're surprised because the odds that 100% of the people in the country would vote the same way are slim – so slim, in fact, that it's more likely that the election was rigged. Indeed, even if 100% of the people *wanted* to vote for the president, it is still unlikely that 100% would vote for the president, given people's surprising capacity to spoil their ballots or vote for the wrong candidate by accident. Of course, total unanimity within a small group is less improbable—and therefore much less surprising—than unanimity in large group.

Let's apply this line of thought to the religious case. Suppose that God exists and that there's publicly available evidence of this fact, but that the strength of the available evidence lies somewhere between inconclusive and overwhelming. Suppose, furthermore, that humans are capable of knowing that God exists *via* this evidence, but that they too are fallible. Sometimes, for instance, they don't take notice of all the evidence available to them; other times, they get distracted by irrelevant information and make bad inferences based upon it. Well, in these circumstances, we can reasonably predict a certain amount of disagreement. Some people will fail to reach the truth. What should we make of total unanimity in such circumstances? This seems like a case in which unanimity is evidence of error. In a world in which fallible reasoners form their judgements based on mixed evidence, total unanimity

would suggest the intervention of some epistemically irrelevant influence on their belief-forming faculties. In short: unanimity is an empirically implausible and normatively untenable goal for anyone who wishes to defend theism by appealing to the wisdom of the crowd.

But even if the first premise were true, the argument is invalid anyway, at least in the form in which it is presented above. For the *consensus gentium* argument to stand a chance, we need to introduce some epistemological principle that can carry us from a premise about what people believe to a conclusion about what exists. Kelly, in his critique of *consensus gentium* arguments (2011), argues that the most promising way to bridge the gap between the premise that everyone believes that theism is true and the conclusion that theism is true is *via* an inference to the best explanation. The crucial premise is roughly: if the best explanation of the consensus that theism is true is that theism is in fact true, then it is epistemically reasonable to believe that theism is true. The principle at work behind this is the following: if the best explanation for some phenomena *F* is that *p*, then given *F*, it is epistemically reasonable to believe that *p*, absent defeaters.

In his response to Kelly, Dobrzeniecki argues that the data point, namely, that a supermajority of the world's population believes that theism is true, is not true in a way that is suitable for use in Kelly's revised *consensus gentium* argument (2018, 136-139). For Dobrzeniecki, the problem is that premise 1 glosses over the huge diversity among people's (in both senses of the word) different beliefs about the nature of the god(s) that they believe in. Mere verbal agreement that God exists – that theism is true – is not sufficient to support the claim that theism is true *via* an inference to the best explanation. After all, he says, when we consider the sheer diversity of the array of beliefs that are bundled together under the banner of theism, isn't the better explanation simply that human beings tend to be religious? In his words, what we've identified as premise 1 is 'defendable only when interpreted as tantamount to the observation that everyone is religious' (2018, 139). For this reason, according to Dobrzeniecki, Kelly's *consensus gentium* argument fails.

Zagzebski (2011) takes a different approach in constructing a *consensus gentium* argument of her own. At the core of Zagzebski's approach is the concept of epistemic self-trust. We're entitled to a presumption in favour of the reliability of our own perceptual and cognitive faculties, unless we acquire some additional reason to doubt the same. An interesting consequence of this plausible view about our pre-reflective epistemic status is the following: since we believe that other people have the same faculties as us, and the same entitlement to the presumption of reliability, then 'there is a presumption in favour of the veridicality of the faculties of other persons until shown otherwise by a further use of [one's] own faculties' (2011, 185). Now, if we have a defeasible entitlement to trust the faculties of others (cf. Foley 2001), then, if they predominantly tend to believe that theism is true, then we're entitled to believe that theism is true *on trust*.

In his response (2018, 145-147), Dobrzeniecki notes that key premise of Zagzebski's argument is the presumption of trustworthiness of other persons' faculties, and the lynchpin of this is the supposed symmetry between their cognitive faculties and one's own. Consequently, Dobrzeniecki argues, the limits of trust in others' faculties are reflected in the limits of our trust in our own faculties. This varies not only with the faculty in question, but also with subject matter and context. If you're *prima facie* entitled to trust your own memory of recent events, then you're *prima facie* entitled to trust someone else's memory reports. And if you're not *prima facie* entitled to trust your own reasoning about conditional probabilities, then you're not *prima facie* entitled to trust someone else's. The question that concerns us is whether the religion is like the first case or the second one: do we have a *prima facie* entitlement to trust our cognitive faculties when it comes to the question of whether theism is true? Dobrzeniecki, following Thomas Aquinas, thinks that the answer is likely to be negative.

It's not our objective to adjudicate this debate between Kelly, Zagzebski, and Dobrzeniecki, although we may comment on their respective positions at certain points in this essay. These are just different arguments based on the same interesting data point and leading towards the same interesting

conclusion. There are others that are beyond the scope of this essay.¹ Rather, we intend to present an alternative defence of *consensus gentium* arguments for theism, appealing to Condorcet’s jury theorem. In the following section, we will introduce Condorcet’s jury theorem. Following that, we will explain how Condorcet’s theorem can help to shore up a defence of theism based on the wisdom of the crowd. We should add the caveat that in this essay we are not going to address the data point that a simple majority of people believe that theism is true in great detail. Of course, this leaves our argument open to the criticism that the first premise is false. That’s fine with us—what interests us more is the matter of bridging the gap between that empirical claim about consensus and the conclusion about what it’s reasonable for us to believe.

§3 Introducing Condorcet’s jury theorem.

Knowledge is widely distributed across society: there have been notable allusions to the ‘wisdom of the crowds’, from ancient times in Aristotle’s *Politics* to the recent rise to prominence of ‘crowdsourcing’ techniques (Surowiecki 2004). Condorcet’s famous jury theorem fleshes out formally some of the insights behind those allusions. The theorem has a non-asymptotic aspect, according to which the majority vote among a group of sincere, competent, and independent people is more likely to be right than are individual people separately, and that probability increases as the number of people does. The asymptotic aspect holds that as the number of those sincere, competent, and independent people approaches infinity, the probability that the majority is correct approaches one. The non-asymptotic part is proven by induction on the number of people, while the asymptotic part is proven by means of the law of large numbers. We will not rehearse the formal proofs here, since others have already done this work (e.g., Ladha 1992). In keeping with the consensus, we take the mathematical correctness of the theorem for granted. Scholars, however, sometimes claim that, although internally consistent, the theorem is not applicable since the three conditions of Sincerity, Competence and Independence are highly unlikely to be jointly satisfied in real life situations.² In this section, we will introduce these conditions and, in section 4, we will consider whether these conditions could be satisfied in the religious case that interests us.

First, however, it’s important to acknowledge a major limitation of the classic Condorcet jury theorem which isn’t relevant in our case. The classic setup concerns a group of people making a choice between only two alternatives by means of majority rule, where exactly one of the two alternatives is correct. Now, many real-world choice situations involve more than just two options. In such cases, Condorcet’s jury theorem quickly becomes untenable, at least without substantial modifications. We, however, are interested in the case regarding the truth of theism. So, not only are we interested in a two-option case, but also, we can safely assume only one of those options is the correct one: theism is either true or false (i.e., some divine person(s) either exists or doesn’t exist). With that caveat in place, we will introduce Competence, Independence, and Sincerity over the following paragraphs.

Competence. If we’re going to base our beliefs about the world on the facts about what other people believe, then we should be concerned about their competence—roughly, how likely it is that their beliefs are true. Here’s how that fits into Condorcet’s jury theorem. Before choosing an option, each person forms a belief regarding which is correct. Each person’s competence is the probability that their belief is correct. Competence requires that the probability that the person’s belief is true is greater

¹ See, for instance, Meierding 1998, O’Briant 1985 and Smith 2020.

² To avoid confusion between technical and everyday uses of these terms, we will capitalise the terms ‘Sincerity’, ‘Competence’, and ‘Independence’ when we are referring to the conditions stipulated by Condorcet’s jury theorem.

than 0.5. That is, for someone to satisfy Competence, they need to be more likely than not to identify the correct alternative. Now, the classical formulation of the theorem states that every person has the same probability of holding the right option. If this were a requirement of Competence, it would certainly make the theorem inapplicable, since Competence would be unlikely to be satisfied. Fortunately, Competence only requires that the competence of each person be above 0.5 (in other words, that their ability to pick the right choice is better than random). But this doesn't require that each person have the same probability, although it is, of course, mathematically convenient to assume that everyone does. Providing each person's competence is better than random and so the mean probability is too, the theorem will still provide us with the desired asymptotic result (even if calculating the mean will be more complex in a real-world situation).³ So some people can be very good at identifying the right answer and others not so good, providing their competence is above 0.5 in order to get the asymptotic result. Having said that, things are slightly more complicated regarding the non-asymptotic result, given that you can have a much more competent minority than the majority. In this case, the non-asymptotic result wouldn't hold. However, if the levels of competence are similarly distributed in both groups, the non-asymptotic result of the theorem would hold.

Independence. If you're surveying a group of people for their opinion, it's important to establish whether they've made up their own mind or they were unduly influenced by external forces. This brings us to Independence, whose role is to ensure that each person provides new, independent information when making a choice instead of simply repeating information with a common source. Otherwise, asking a group of people for their opinion would be like reading different token copies of the same newspaper. Ideally, when we survey groups for their opinions, we want them to be more like different sources of information. Imagine that all the people except for one choose an option by *blindly* following that one person. That would certainly not be a case of 'wisdom of the crowds' but maybe, depending on that one's person competence, a case of 'wisdom of the one'. Just like different token newspaper copies, each of those people would be merely echoing the one person's judgement. It is standardly thought that Independence requires that the beliefs of all the people be statistically independent. Just like the tossing of two fair coins are independent (knowing the result of one coin toss doesn't increase one's chances of correctly guessing the result of the other one), two people's chances of getting the right answer are independent if neither gets a higher chance of getting the right answer given that the other one does. This doesn't rule out interaction between people (that is, it doesn't require

³ To appreciate this informally, consider an extension of Estlund's (2008, 224) coin example. He first asks us to consider flipping a fair coin. If you toss it just a few times (say, 10 times), it is not at all unlikely that an unequal number of heads and tails results would happen (say, 70/30). But if you do it many times (say, 1,000 times), then it is much more unlikely that such unequal outcome will take place. The more times you toss it, the more likely you will reach the 50/50 outcome. Now, imagine you have a coin which has 51% chance of being heads when tossed. In this case too, you can expect to get closer to the 51% heads outcome the more you toss it. Next, instead of tossing the one unfair coin many times, think of tossing many unfair coins whose chance to be heads is 51%. In this case too, you would expect the same result as before: the bigger the number of unfair coins, the bigger the chances that the frequency of heads would come to 51%. Now, change coins for people and tails and heads for the wrong and right answer and you can informally appreciate the logic behind the Condorcet jury theorem. If the people have 51% chances of delivering the right answer, the more such people you have, the more likely you are to get a majority of correct answers. Moreover, when that number goes to infinity, the chances of the right answer being delivered by the majority converges to 1. And if you think of those people as having different chances (although bigger than random) of delivering the right answer, the same will be true.

causal independence) but restricts positive correlations of their decisions. Independence is a claim about the relation between those decisions and not the idiosyncratic influences each person may have. In other words, there can be some deference and deliberation among the people. After all, interpersonal discussions don't imply consensus. In fact, just as they can make decisions *more* interdependent (say, by generating homogenization and so causing positive correlations (Luskin et al. 2022)), they can also make decisions *less* interdependent (say, by allowing people to realize that they are paying too much attention to one source of information or motivating them to make up their own minds, thereby increasing the independence of the decisions). Certainly, if some common cause *systematically* makes them take the same decisions, then they wouldn't be independent. But that doesn't imply that all common causes are problematic. In fact, it is epistemically desirable that the state of the world related to the alternatives considered be a common cause. After all, we want people to be competent: to be able to track the true state of the world. Independence is understood as statistical independence given the state of the world. Moreover, given that it is *via* (non-misleading) evidence that we are acquainted with the (true) state of the world, this then requires that Independence be conditionalized on such evidence as well as Competence (Dietrich and List 2004).⁴ Having said that, other possible common causes, such as shared ideology and systematic cognitive shortcomings (or the one exemplified in the case of unquestioning deference on a person), violate Independence. It makes people take the same decisions (even without direct interaction in some cases), hence them being positively correlated.

Having noted the above connection between Competence and Independence, it's worth pointing out that not all distal and proximate causes of belief will affect them equally. For example, both idiosyncratic and systematic cognitive shortcomings will have a negative effect on the person's competence but only the latter is guaranteed to have a negative effect on the decisions' independence. Different cognitive shortcomings, like different opinion leaders with diverse views, can cancel each other out. So, if some people are subject to different common causes, Independence could still be met by all the people involved in the decision. That is, some non-extreme cases of dependence (say, two sub-groups of people following blindly opinion leaders with contrary views) can be tolerated, but extreme cases (such as everyone blindly following the one person's opinion) cannot. Now, while non-extreme cases are probably quite common multicultural and diverse societies like ours, it is less likely that the extreme ones are. Furthermore, whether the non-extreme cases do or do not affect Independence is of course an open empirical question. There certainly is the possibility that they do, but it should be clear that whether a particular case does so cannot be established purely from the armchair.

Sincerity. Of course, it's one thing to establish whether people are competent and independent: it's quite another to establish whether they're honest. It doesn't matter how competent someone is, if you have reason to believe that they might be trying to deceive you. This fact, often discussed in the epistemology of testimony, is relevant here too. For Sincerity to be satisfied, it must be the case that, when deciding which alternative to opt for, members of the voting population do so sincerely. The option they choose (e.g., by voting) is the option they believe is correct. This is an important requirement since it allows us to bridge the gap between what people say they believe and what they actually believe. It is only under these circumstances that we are justified in applying Competence and Independence in our aggregation of the judgements of the crowd. If a person is insincere, then even if they're competent and independent, we cannot take what they say as reflecting their beliefs.

⁴ As Dietrich and List put it, the decisions 'are independent from each other not once the state of the world is given, but once the evidence is given' (2004, 242). Moreover, given that Competence is conditional on the evidence, Dietrich and List speak of tracking the 'ideal observer' (instead of the state of the world), who 'gives the ideal (best possible) interpretation of the available evidence' (243).

§4 Condorcet, *consensus gentium*, and theism.

In this section, we'll explore a number of ways in which the insights from Condorcet's jury theorem show that *consensus gentium* arguments are not so easy to dismiss after all, and that there are conditions under which they would lend support for theism. Since our goal is to present a defence of *consensus gentium* arguments for theism, we will argue that, for all we know, Competence, Independence, and Sincerity are jointly satisfied when it comes to the population's views about theism. We do this by presenting a presumptive case for the satisfaction of each condition of Condorcet's jury theorem. We will refrain from defending the stronger position according to which the *consensus gentium* argument shows that it is epistemically reasonable to believe that theism is true. A brief clarification: there is a weakness in the analogy between Condorcet's jury theorem and *consensus gentium* arguments. In Condorcet's jury theorem, public opinion is surveyed *via* a voting process. In conventional *consensus gentium* arguments, there is no such voting process. Rather, reports of belief are taken to be the equivalent of votes on the truth of theism.

Competence. Matheson (forthcoming) expresses the worry that it is hard to know whether Competence is satisfied, and what's more, that the evidence suggests that it is not. He points out that it's difficult to determine how reliable we are regarding the truth of theism without begging the question. To try to overcome this, he suggests that our reliability when it comes to our other religious beliefs could provide an indication of how reliable we are with respect to our beliefs about theism. Matheson's concern is that the widespread disagreement between religious believers regarding all manner of religious issues (besides the existence of God) is higher-order evidence that we're not competent epistemic agents with respect to (a) those topics and, by extension, (b) theism. In the following paragraphs we'll suggest some ways to avoid such concerns. First, it's not clear to us that evidence of disagreement about the finer details of religion is evidence of the incompetence of judgements about theism. Consider this example: particle physicists agree that subatomic particles exist, but they disagree about how they interact with each other. But their disagreement regarding the nature of subatomic particles doesn't give us reason to believe that the scientists' belief that subatomic particles exist is unreasonable. Indeed, it is interesting that the consensus on theism appears to be the exception: perhaps that's evidence that humans are competent when it comes to the broad question of theism, but incompetent when it comes to narrower questions of religious doctrine. In any case, for Condorcet's jury theorem to work, the voting group must be minimally epistemically competent. In probabilistic terms, this means that the average probability that a member of the voting group reaches a correct judgement with respect to the issue at hand must be greater than 0.5. Our question, then, is the following: how could individuals and groups satisfy Competence in a way that is relevant to the matter of theism?

Let's start at the extremes. First, let's consider the case in which all members of the population are highly epistemically competent, especially when it comes to matters of theism. Of course, under these circumstances, Competence is likely to be satisfied. This case needn't concern us much. Assuming that Independence and Sincerity are satisfied, if the majority (50%+1) of a sufficiently large, highly competent group votes that p , then it is epistemically reasonable for us to believe that p based on their beliefs. As it happens, it's unlikely that the average person is a highly competent epistemic agent. That's not to pour scorn on the intellectual abilities of ordinary people—rather, the point is that while most people may count as (minimally) competent, far fewer qualify as highly competent experts, and less so when it comes to the thorny issue of the truth or falsity of theism. Such optimism about the average competence of the population would weaken our defence of the *consensus gentium*.

Now let's turn to the other extreme: suppose that most people are incompetent epistemic agents when it comes to the matter at hand. In probabilistic terms, we might say that they have a probability below 0.5 of reaching the correct judgement about the truth or falsity of theism. It goes without saying that, under these circumstances, the fact that a majority of the group believes that theism is true does

not grant us good epistemic reasons to believe that it is, in fact, true. In fact, the reverse theorem results are to be expected if the population is incompetent, as just described. In this case, the logic of the jury theorem allows us to learn about the world that what the majority holds is most likely false and the bigger the number of people involved the bigger the chances of it being so. So, the fact that the consensus is that p gives us a reason to believe that not p !

An interesting and, for our purposes, relevant insight from Condorcet's jury theorem is that minimal competence (providing Independence and Sincerity are satisfied) is sufficient for delivering the asymptotic and non-asymptotic results (just like minimal incompetence is sufficient to deliver the reverse asymptotic and non-asymptotic results). Does the population meet this minimal demand for competence?⁵

A plausible story about how people come to hold their belief in theism is that their beliefs are a result of their upbringing. The same may be said of people who do not believe that theism is true. The thought, then, is that most people believe in theism because they were raised in religious families and brought up in religious communities. But that doesn't mean that they cannot satisfy Competence.⁶ Let's suppose, for argument's sake, that these beliefs are transmitted from one generation to the next *via* testimony. If most people's belief in theism is based on testimony, then we might wonder whether it's reasonable for them to hold the religious beliefs that they learn as children. This leads us back to the issues of trust introduced by Zagzebski and discussed by Dobrzeniecki. The difference between our view and Zagzebski's, and what makes our view immune from Dobrzeniecki's criticism, is that Condorcet's jury theorem doesn't require people to be highly competent. So long as they're minimally competent, the law of large numbers does the rest of the heavy lifting. With a sufficiently large population of minimally competent, independent, and sincere people, the probability of reaching the truth converges on 1.0.⁷ (As we can see here, Competence and Independence are practically intertwined but can be separated conceptually.)

Having said that, how competent people actually are in the matter at hand is partly an empirical issue that lies outside the scope of a philosophy paper. Crucially, one cannot wholesale lump believers or non-believers into one group or the other without considerable evidence. So, without the appropriate empirical work, we cannot, for example, assume that characteristically non-believers are (significantly) more competent than believers. Moreover, for the sake of charity, we can assume that levels of

⁵ One way to establish that theists (and atheists alike!) satisfy Competence is by appealing to the notion of humankind's supposed *sensus divinitatus* (see Plantinga 2000 for perhaps the most influential contemporary treatment of this topic). One question we might ask about the *sensus divinitatus* is, how reliable a cognitive faculty ought to be? Or to put it another way, how bad could a good-enough *sensus divinitatus* be? Let's suppose that the *sensus divinitatus* has a probability of functioning correctly that is ever so slightly greater than 0.5. In other words, that chance that a person's *sensus divinitatus* is functioning correctly is only slightly better than flipping a coin. That might not be sufficient for getting individual believers to a state of knowledge with respect to the existence of God, but it is sufficient to satisfy the Competence condition of Condorcet's jury theorem. And, on the basis of the relevant *consensus gentium* argument, the individual believers may know that some such divinity exists.

⁶ This is related to a debate in epistemology about the rationality of 'nurtured beliefs', that is, the beliefs that one holds because of one's upbringing. The worry is that we seem to hold many of our most important beliefs just because we were raised that way, and that doesn't seem to offer those beliefs the kind of epistemic support that is required for epistemic rationality. For key contributions to this debate, see Bogardus 2013, Cohen 2000, Elga 2007, Schoenfield 2014, Vavova 2018 and White 2010.

⁷ Although as footnote 5 suggests, things are not as straightforward as this, given that Competence is conditional on the available evidence and so the probability of reaching the truth will depend on it.

competence are similarly distributed in both groups, in which case the non-asymptotic result of the theorem would hold. Finally, there is nothing in the theorem itself that requires generalizing this competence to all issues. What matters for Competence is the capacity of the people to choose the right answer within the topic considered.

Independence. How dependent can a sufficiently independent member of the voting group be? Again, let's start at the extremes: at one extreme, all people would be unquestioningly, dogmatically deferent to one opinion leader. At the other extreme, all people would be completely independent. In between the extremes, there are more interesting cases. To illustrate this, imagine a world in which there are hundreds of thousands of sects. Each sect is led by a minimally competent, independent, opinion leader *and* all the members are completely deferential to the leader of their respective sects. Counter-intuitively, this is enough to satisfy Independence, although reducing the number of actual independent members to be taken into account (from billions to hundreds of thousands).

We don't, however, live in either of these kinds of world. We live in a multicultural society with a diverse panel of opinion leaders and some people might just not follow them. Now, the people who follow opinion leaders (on both sides) can cancel each other out. Whether that is the case it is, of course, an empirical matter but, providing a sufficiently big number of people aren't followers, then Independence can still be satisfied. Moreover, it is unlikely that most followers (on both sides) blindly defer to their opinion leaders. Even within hierarchical religious organizations, such as the Anglican Church, there is considerable room for critical thinking and independent thought.⁸ And, importantly, any person who responsibly accepts the testimony of an opinion leader wouldn't blindly defer to her. To repeat, Independence doesn't require causal independence and interaction between the people is allowed.

However, as mentioned above, other common causes, other than common opinion leaders who are slavishly deferred to, can violate Independence. A shared ideology is one such example. But again, just like in the case above, in multicultural and diverse societies we would expect there to be several ideologies and, more importantly, we would expect only some people to be committed to them. So, these different ideologies can cancel each other out and, significantly, we can expect the commitment of some people to their ideologies not to be blind. Similarly, a systematic cognitive shortcoming could count as common cause that violates Independence. For example, a prejudice or bias that affects everyone.⁹ But, again, we can expect there not be some such single common cause for everyone. For all the similar features that we *qua* human beings have, it is also clear that different people will be subject to different prejudices and biases. This is in part because different people are likely to have different evidence regarding the relevant state of the world and so different prejudices and biases are likely to be triggered. Moreover, and significantly, each of these prejudices and biases that actually affect people are likely to do it in different ways; say, the tendency to ignore evidence contrary to what

⁸ Interestingly, it has been argued from an evolutionary point of view that people benefit from being listened to, which creates pressure for original opinions within the group and so for people to be independent-minded (de Courson et al. 2021). This then suggests that at least some people are likely not to blindly follow others.

⁹ Recent work in the cognitive science of religion has attempted to show that there are common causes for religious belief in humans which are rooted in the mechanics of our psychology rather than in our judgment about the world. If such an approach were to yield a successful debunking explanation of religious belief, that would of course be a problem for *consensus gentium* arguments in favour of theism. Interestingly, common causes for believing in atheism have been sought by people defending theistic worldviews, such as Kelly James Clark (2019).

one believes (i.e., the confirmation bias) can be stronger in one person and weaker in another one. So, again, it is not straightforwardly clear that people don't satisfy Independence.

Sincerity. Yet again, the relevant question here is how insincere can a sufficiently sincere member of the group be? Let's return to the extremes. At one extreme, all the members of the voting population would be utterly sincere when it comes to matters of theism. At the other extreme, all the members of the group would be utterly insincere with respect to their beliefs about theism. It is obvious that if the first scenario obtains, then the Sincerity condition of Condorcet's jury theorem would be satisfied. Moreover, it's obvious that if the second scenario obtains, then Sincerity is not satisfied. The challenge arises when we try to determine where we should draw the line.

Nevertheless, it is probably safe to assume the Sincerity requirement is typically satisfied in the case at hand. It is, for example, very likely that most religious believers sincerely hold their religious views. It is, of course, a possibility that some people, for cultural or political reasons, a person claims to be (say) a Christian, even though they do not believe that theism is true, among other things. The existence of such cases is only problematic for Condorcet's jury theorem if they are widespread. Likewise, some people might make the strategic decision, *à la* Pascal, to believe in the existence of God. Pascal famously argued that we have prudential reasons for believing in God. A common rebuttal to this position is that, given the doxastic voluntarism is false, we cannot believe at will and therefore taking Pascal's wager is psychologically impossible. Pascal famously argued that the seeking non-believer should join a religious community and act as if they believed, and that this will bring about religious belief.¹⁰ Is Pascal's libertine a sincere believer? It depends on your reading of *Pensées*. If you think it's an economic argument, then the libertine is probably insincere. If you think that Pascal's purpose is to provide a psychological lifeline for someone whose would-be belief is plagued by doubts, then you might have a more sympathetic view of the sincerity of Pascal's seeking non-believer. However, having noted the above, it is probably safe to assume that Sincerity mostly holds among the religious and non-religious believers at least concerning the case at hand, although this is a contingent matter. If so, we can assume that Competence and Independence apply to the actual choices and so it is not fruitless to examine whether the latter requirements can be met.

Anyway, it's worth pointing out that the position we're defending here can withstand a small number of insincere voters. The important thing is that if a big enough number of them is sincere, then we can still expect the asymptotic and non-asymptotic results of Condorcet's jury theorem. And to repeat, we can expect that, with regard to the matter at hand, most people are in fact sincere. Of course, whether they *actually are* is an open empirical question. There certainly is the possibility that they aren't, but that (just like the competence and independence of people) cannot be established purely from the armchair.

§5 Conclusion.

We don't know whether *consensus gentium* arguments actually offer support for theism. But we're more confident that they're interesting arguments, and that they are worthy of more attention than they have hitherto received in the philosophy of religion. On more than one occasion we have set potential problems aside on the grounds that they're empirical matters.¹¹ While these potential problems aren't

¹⁰ Garber (2009) has argued that Pascal's regimen could not yield epistemically rational religious belief. Responding to Garber, Efir and Warman (2021) have argued that, for all we know, Pascal's regimen could yield epistemically reasonable religious belief.

¹¹ Of course, whether any of these conditions are satisfied is a contingent matter. We have made the case that they could be, at least in a society that resembles ours. It's an interesting thought that there are

purely empirical, our aim is to show that at the theoretical level, *consensus gentium* deserve more credit than they have been given. We hope that, if we have convinced the reader of one thing, it is that whatever the weaknesses of the *consensus gentium* argument may be, their probative value is nevertheless underappreciated, and they have been unfairly maligned as a consequence.

References.

- Bogardus, Tomas (2013). The problem of contingency for religious belief. *Faith and Philosophy* 30 (4): 371-392.
- Clark, K. J. (2019). *God and the Brain: The rationality of Belief*. Eerdmans Publishing company, Michigan.
- Cohen, G. A. (2000). *If you're an egalitarian, how come you're so rich?* Cambridge, MA: Harvard University Press.
- de Courson, B., Fitouchi, L., Bouchaud, J-P., & Benzaquen, M. (2021). "Cultural Diversity and Wisdom of Crowds are Mutually Beneficial and Evolutionary Stable". *Scientific Reports* 11, 16566.
- Dietrich, F. and List, C. (2004). "A Model of Jury Decisions Where All Jurors Have the Same Evidence", *Synthese*, 142:175-202.
- Dobrzeńiecki, Marek (2018). Is the fact that other people believe in God a reason to believe? Remarks on the consensus gentium argument. *European Journal for Philosophy of Religion* 10 (3): 133-53.
- Efird, David and Warman, Jack (2021). After Pascal's wager: On religious belief, regulated and rationally held. *International Journal for the Philosophy of Religion*. 90 (1): 61-78.
- Elga, Adam (2007). Reflection and disagreement. *Noûs* 41 (3):478–502.
- Estlund, D. (2008). *Democratic Authority: A Philosophical Framework*. Princeton: Princeton University Press.
- Foley, R. (2001). *Intellectual Trust in Oneself and Others*. Cambridge: Cambridge University Press.
- Garber, D. (2009). *What happens after Pascal's wager?* Marquette University Press.
- Gunn, L. J., Chapeau-Blondeau, F., McDonnell, M. D., Davis, B. R., Allison, A. & Abbott, D. Too good to be true: When overwhelming evidence fails to convince. *Proceedings of the Royal Society A* 472: 20150748.
- Joyce, George Hayward (1923). *Principles of Natural Theology*. New York: Longmans, Green & Co.
- Kelly, Thomas (2011). *Consensus gentium: Reflections on the 'common consent' argument for the existence of God*. In Clark & Van Arragon (eds.), *Evidence and Religious Belief*. Oxford: Oxford University Press.
- Ladha, K. (1992). "The Condorcet Jury Theorem, Free Speech and Correlated Votes", *American Journal of Political Science*, 36:617-634.
- Luskin, R., Sood, G., Fishkin, J. and Hahn, K. (2022) Deliberative Distortions? Homogenization, Polarization, and Domination in Small Group Discussions. *British Journal of Political Science* 52, 1205-1225.

possible worlds in which Competence, Independence, and Sincerity are all satisfied 'with flying colours', so to speak. In such worlds, the probative value of *consensus gentium* arguments would be considerably greater. Conversely, and as the objections we have considered suggest, in a different possible world where incompetence, insincerity, and herd mentality are rife, *consensus gentium* arguments would have little probative value at all.

- Matheson, Jonathan (forthcoming). The common consent argument. In Ruloff (ed.), *Contemporary Arguments in Natural Theology*. London: Bloomsbury.
- Meierding, Loren (1998). The consensus gentium argument. *Faith and Philosophy* 15 (3):271-297.
- O'Briant, Walter H. (1985). Is there an argument *consensus gentium*? *International Journal for the Philosophy of Religion* 18: 73-79.
- Plantina, Alvin (2000). *Warranted Christian Belief*. New York: Oxford University Press.
- Schoenfield, Miriam (2014). Permission to believe: Why permissivism is true and what it tells us about irrelevant influences on belief. *Noûs* 48 (2):193-218.
- Surowiecki, J. (2004). *The Wisdom of the Crowds*. London: Abacus.
- Smith, Tiddy (2020). The Common Consent Argument for The Existence of Nature Spirits. *Australasian Journal of Philosophy* 98 (2): 334-348.
- Vavova, Katia (2018). Irrelevant influences. *Philosophy and Phenomenological Research* 96 (1): 134-152.
- White, R. (2010). You just believe that because... *Philosophical Perspectives*, 24: 573-615
- Zagzebski, Linda (2011). Epistemic self-trust and the *consensus gentium* argument. In Clark & Van Arragon (eds.), *Evidence and Religious Belief*. Oxford: Oxford University Press.