HOW SKEPTICAL IS THE EQUAL WEIGHT VIEW?

BRANDON CAREY AND JONATHAN MATHESON

1. Introduction

Much of the literature on the epistemology of disagreement focuses on the rational responses to disagreement, and to disagreement with an epistemic peer in particular. The Equal Weight View claims that in cases of peer disagreement each dissenting peer opinion is to be given equal weight and, in a case of two opposing equally-weighted opinions, each party should adopt the attitude which 'splits the difference'. The Equal Weight View has been taken by both its critics and its proponents to have quite drastic skeptical ramifications given contingent empirical facts that we are aware of regarding disagreements in philosophy, religion, science, and politics. In this paper, we begin by clarifying the central claims of the Equal Weight View (Section 2) and then examine two routes from the Equal Weight View to skepticism about such matters that have been explored in the literature. The first claims that our awareness of peers or experts who disagree with us about such issues requires that we abandon our beliefs on these issues (Section 3). The second claims that our awareness of merely possible peers or experts who disagree with us requires us to abandon our beliefs (Section 4). We find both routes from the Equal Weight View to a form of skepticism defective. However, there are nearby considerations, explored in Sections 5 and 6, which (for better or worse) do lead to at least some skeptical consequences for the Equal Weight View, albeit for different reasons.

2. What is the Equal Weight View?

According to the Equal Weight View (hereafter EWV), in cases of peer disagreement you ought to give your peer's conclusion equal weight unless you have an independent reason for discounting your peer's opinion.[1] In giving your peer's conclusion equal weight you recognize that your opinion is no more likely to be correct (or incorrect for that matter) than your peer's *simply because it is your own*. Epistemic peers regarding p are in an equally good epistemic position regarding p, and so each is as likely to be correct about p as the other. So, giving each peer opinion equal weight seems like the reasonable thing to do. EWV can be motivated by the claim that the weight that should be given to an opinion regarding p corresponds to the epistemic position with regard to p that the person who has the opinion regarding p is in. The stronger the epistemic position regarding p—the more likely you are to be correct about p—the more weight your opinion should be given. So, equality of epistemic position regarding p will call for equality of weight regarding the opinions regarding p.

While the prescription to give a peer's opinion equal weight does rule out *some* ways of responding to a case of peer disagreement, it does not narrow down the rational responses to disagreement to a unique answer. For instance, giving both your opinion and your peer's opinion *no weight* is one way of giving both of your opinions equal weight. That said, EWV is typically understood as claiming that not only should you give the conflicting opinions equal weight, but also that you should give them *substantial* weight, at least in a case of disagreement between reliable evaluators of the evidence.[2] In cases of two-person disagreement like this, EWV has been taken to require "splitting the difference"—adopting the doxastic attitude halfway between the

---

[1] For more on what it takes for such a reason to be independent, see Christensen (2007).

[2] For more on the need for this qualification, see Matheson (forthcoming).

conflicting peer attitudes.[3] So, if S1 and S2 are both parties to a peer disagreement where S1 believes p to degree 0.8 and S2 believes p to degree 0.2, then having discovered the disagreement, S1 and S2 would 'split the difference' by each believing p to degree 0.5. The 'splitting the difference' application of EWV fits naturally with a fine-grained doxastic picture. On a doxastic picture where one's doxastic options are limited to three: all-out-belief, all-out-disbelief, and all-out-suspension of judgment, "splitting the difference" becomes more complicated. While it is apparent how a disagreement between an individual that all-out-believes p and an individual that all-out-disbelieves p is to be handled (each should split the difference and suspend judgment), it is hard to see how a splitting the difference maneuver could even be accomplished in a disagreement between an individual that all-out-believes p and an individual who suspends regarding p.[4]

In what follows we will be understanding EWV as claiming the following three things:

(i)   If S1 and S2 are aware that they are epistemic peers regarding p, and there is no independent reason for discounting either opinion about p, then S1's opinion and S2's opinion about p should be given equal epistemic weight.

(ii)   If S1 and S2 are aware that S1 has a better epistemic position toward p than S2, and there is no independent reason for discounting S1's opinion, then S1's opinion regarding p should be given more weight than S2's.

---

[3] As we will see, things are more complicated when there are more than two disagreeing parties.

[4] See Kelly (2010, 117) for more on this problem.

(iii)   If S1's opinion and S2's opinion about *p* are to be given equal weight, then if S1 and S2 are aware that they disagree regarding *p*, and have no independent reason to discount either opinion regarding *p*, they ought to each 'split the difference'.[5]

While this understanding of EWV builds more into the view than merely giving each peer opinion equal weight, we believe that this richer claim is what both proponents of EWV and its opponents have had in their sights.

To better see how EWV is to be applied to cases of disagreement, it will be helpful to distinguish between two kinds of evidence: first-order evidence and higher-order evidence.[6] *First-order evidence* regarding *p* is evidence directly bearing on *p*. For instance, there appearing to be a red ball in front of me is first-order evidence that there is a red ball in front of me. In contrast, *higher-order evidence* regarding *p* is evidence about the first-order evidence regarding *p*. For instance, evidence about the reliability of my perceptual experiences is evidence that the first-order evidence of there appearing to be a ball in front of me is in fact good evidence that there is a ball in front of me. A species of higher-order evidence, one directly pertaining to issues of disagreement, is the opinions of reliable evaluators of the evidence. That a reliable evaluator of the evidence has evaluated a body of evidence to support *p* is higher-order evidence that that body of evidence in fact supports *p*. Since reliable evaluators of the evidence typically evaluate bodies of evidence correctly—that is, since they typically adopt the doxastic attitude toward a proposition that fits the evidence—evidence that such an individual has judged a body of

---

[5] While these conditions have each been formulated in terms of what S1 and S2 are aware of, we think that it is also plausible that analogous principles hold where the fallible "justified in believing" is substituted for the factive "aware that."

[6] This distinction is not new to the literature on disagreement: see Kelly (2005) and Matheson (2009).

evidence to support *p* is higher-order evidence that the body of evidence in question does in fact support *p*.

It is important to note that, according to EWV, what any person should believe in a case of disagreement is *entirely* a matter of the higher-order evidence—it entirely depends on that individual's evidence about the opinions of rational individuals on the matter.[7] Notably, the opinions of the experts on the matter will be weighted most heavily since those are the opinions of the people in the strongest epistemic position on the matter. On this picture, to find out what I should believe about disputed proposition *p*, I look at what others believe about *p*, and weigh their opinions according to the epistemic position of the individuals holding the opinions with respect to *p*.

In the next two sections, we will investigate several ways in which it might be erroneously thought that EWV has skeptical consequences. In Sections 5 and 6, we will make a positive case that EWV indeed has at least some skeptical consequences, though not quite for the reasons it has been thought to.

3. Is EWV Skeptical?—The Case from Actual Disagreement

EWV has been taken both by its critics and by its proponents to have skeptical ramifications, at least regarding many of our political, religious, ethical, and scientific beliefs, given contingent empirical facts regarding disagreements that we are aware of in these fields. Thus, Richard Feldman (2007, 216) claims:

---

[7] For a criticism of EWV on this point, see Kelly (2010).

[EWV] is also a skeptical view, in the limited sense that it denies the existence of reasonable

beliefs in a significant range of cases. This may seem to be a distressing conclusion.

In Feldman (2006, 217), he reinforces this idea:

[O]n many [issues in philosophy, religion, science, and public policy] about which you have a

belief, informed and intelligent people disagree with you. The question I will raise concerns the

reasonableness of maintaining your point of view in the light of such disagreements. My

conclusion will be that, more often than we might have thought, suspension of judgment is the

epistemically proper attitude. It follows that in such cases we lack reasonable belief and so, at

least on standard conceptions, knowledge. This is a kind of contingent real-world skepticism that

has not received the attention it deserves.

David Christensen (2009, 757-758) echoes this sentiment:

The most obvious motivation for [opposing] views on disagreement flows from the degree of

skepticism that [EWV] would seem to entail. There must be something wrong, the thought goes,

with a view that would counsel such widespread withholding of belief. If you have an opinion on,

for example, compatibilism about free will, scientific realism, or contextualism about knowledge,

you must be aware that there are very intelligent and well-informed people on the other side. Yet

many are quite averse to thinking that they should be agnostic about all such matters. The aversion

may be even stronger when we focus on our opinions about politics, economics, or religion.

The form of the argument from EWV and contingent facts about disagreement to skepticism

about many of our beliefs about controversial propositions can be reconstructed as follows,

letting C be the class of "controversial" propositions in politics, religion, ethics, science, etc. and

S be some subject:

6

(1) For each proposition p in C, S is aware of an epistemic peer or superior whose opinion S has no independent reason to discount with whom S disagrees about p.

(2) If (1), then if EWV is correct, then for each proposition p in C, there is an individual with whom S should at least split the difference with regard to p.

(3) If EWV is correct, then for each proposition p in C, there is an individual with whom S should at least split the difference with regarding p. (1, 2)

(4) If (3), then if EWV is correct, then for each proposition p in C, it is irrational for S to believe p.

(5) If EWV is correct, then for each proposition p in C, it is irrational for S to believe p. (3, 4)

Call this argument the 'Skeptical Consequences Argument' or SCA. If we satisfy the SCA—if the premises are true when each of us is used as the value of S—then it is irrational for us to believe any controversial proposition in politics, religion, ethics, and science. Clearly, if it is irrational for us to believe any of these propositions, we also do not know them, and so skepticism follows.

Do we satisfy the SCA? It is plausible that we satisfy premise (1), since for every proposition p in C we are aware of at least one individual who is in at least as good an epistemic position toward p as we are, yet disagrees with us regarding p. This is at least part of the explanation as to why these propositions are controversial. Further, it doesn't initially seem that we have any special reason to discount this disagreeing individual's opinion regarding p, especially one that is independent of the disagreement itself. According to claims (i)–(iii) of EWV, we also satisfy premise (2): if we are aware of a disagreeing peer whose opinion we have no independent reason to doubt, by (i) we should give that peer's opinion equal weight, and so by (iii) we should split the difference, whereas if we are aware of a disagreeing superior whose opinion we have no reason to doubt, by (ii) we should give that superior's opinion more weight than our own and thus presumably make an even more dramatic change in our opinion regarding

p. (3) simply follows from (1) and (2). It is helpful to think about premise (4) on a tripartite doxastic taxonomy. On this doxastic picture, if we are required to *at least* split the difference with this disagreeing individual, then this will minimally call for us to abandon our belief about *p*. Hence, it will be irrational for us to believe what we did about *p*.

While the SCA is appealing, there are good reasons to think that most of us in fact do not satisfy it. First, premise (4) looks implausible for many subjects and propositions on a fine-grained doxastic taxonomy. Many instances of 'splitting the difference' will result in a non-skeptical doxastic attitude toward *p*. For instance, if S1 believes *p* to degree 0.9 and S2 believes *p* to degree 0.8, then splitting the difference has each believing *p* to degree 0.85. Such an attitude toward *p* is far from a skeptical one. So, the path from splitting the difference to skepticism is neither straight nor clear. Often, the attitude which gives each disagreeing peer opinion equal weight and gives the opinions of those in better epistemic positions more weight will not be a skeptical attitude.

That said, skeptical worries might re-emerge when we consider the fact that regarding a great deal of the propositions in C it will be the case that there is an epistemic peer or superior who *significantly* disagrees with us about *p*—an individual who adopts a drastically different doxastic attitude toward *p*. This is plausible since all the propositions in C are controversial, and controversy does not typically arise from such small disagreements as whether one ought to believe *p* to degree 0.8 or 0.9. Given this empirical fact, it may seem like the above objection can be handled, since giving the opinion of the significant disagreer the weight it deserves will require significant doxastic conciliation on our part, moving us away from any doxastic attitude that affirms *p*.

Whether or not such considerations can make it plausible that we satisfy (4) for every proposition in C, we think that there is a serious problem with (1). While it is true that we often

encounter conflicting opinions, it is also true that we often have an independent reason to discount these conflicting opinions. This reason comes from what we know about the distribution of other opinions on the matter. EWV is not committed to the view that a *sole* peer dissenter regarding p renders your belief that p unjustified.[8] Even if one of your epistemic peers disbelieves that $1 + 1 = 2$, you have at least one good reason to think that he is mistaken—a reason that is independent from your reasoning about the matter. This reason comes from what you know about other epistemic peers and superiors, and what they believe about whether $1 + 1 = 2$. If each peer opinion is to be given *equal weight*, then a *sole* dissenting peer opinion need not render your doxastic attitude toward p unjustified. The weight to be given to this disagreeing opinion is swamped by the weight to be given to the other agreeing opinions among your other epistemic peers and superiors.

Similar considerations apply when the minority dissenter is an epistemic superior. Graham Priest is no doubt our epistemic superior regarding claims of logic, and we are aware that he significantly disagrees with us on at least one such matter in holding that contradictions can be true. Does EWV mandate that we adjust our attitudes significantly on the matter of whether contradictions can be true? It seems not, since we are also aware of widespread agreement regarding the impossibility of true contradictions among other epistemic peers and superiors of ours. We have reason to discount Priest's opinion on the matter, and this reason is independent from our own reasoning about whether contradictions can be true.[9]

---

[8] For an argument that it does, see Lackey (2010).

[9] Note that it does not follow from these considerations that Priest is irrational in maintaining his opinion, despite his awareness of the same widespread agreement that no contradiction can be true. He may have other reasons that we lack for discounting the opinions of those who disagree with him.

Following this train of thought, it can seem that a skeptical attitude is hardly ever called for (if at all) by EWV. Take any widespread disagreement that we are aware of regarding one of our beliefs in a proposition in C. Peer opinions (and expert opinions) on the matter are divided, no doubt, but they are not divided in a *perfectly symmetrical way*. That is, there will typically be *some* reason to think that one side of the dispute is mistaken, coming from the fact that it is the minority view among the relevant peers and experts. According to EWV, what we should believe about any disputed matter is the result of the higher-order evidence, particularly the evidence about peer and expert opinions on the matter. As such, what we should believe is the result of something like an 'epistemic election'. The election doesn't give everyone a vote (for instance, small children are left out), and it has it that some votes count more than others (for instance, the relevant experts' opinions count for more). The better one's epistemic position on the matter, the more weight one's vote gets. In such an 'epistemic election' the result is very rarely a tie, and so there will almost always be *some* reason to privilege one of the disputants' sides. So, it does not seem plausible that we satisfy (1) with respect to very many propositions. Rather, it seems that on neither doxastic picture does EWV readily lead to skepticism about many of our beliefs in propositions in C, even when EWV is coupled with contingent facts about disagreements regarding these matters. The SCA fails.

That said, the considerations brought forward in rejecting instances of (1) have the consequence that EWV and contingent facts about disagreement will often avoid skeptical consequences by mandating that you disbelieve a proposition that you went into the disagreement believing. That is, while the 'epistemic election' will often not end in a tie calling for a skeptical attitude, it will often not come out as you would like, calling for you to adopt quite a different non-skeptical attitude toward *p* than you previously had. Some may find this consequence even *worse* than a skeptical consequence, since in some sense the loss is greater. Not only were you

required to abandon your belief that *p*, you were required to take up a belief that not-*p*! That said, such a non-skeptical consequence does permit one to rationally hold beliefs concerning propositions in C (even if not the ones you wanted), and getting to have some rational beliefs on these matters might be in some sense better than being forced to have skeptical attitudes toward all such propositions.

4. Is EWV Skeptical?—The Case from Merely Possible Disagreement

Another reason that skepticism might be thought to follow from EWV comes from thinking about merely possible disagreers.[10] Since we almost always have independent reasons to doubt either our own opinions or the opinions of those who disagree with us, the antecedents of (i) and (iii) are almost never satisfied. However, it is (metaphysically) *possible* that these conditions be satisfied—we *could* encounter a peer whose contrary opinion we genuinely have no independent reason to doubt (while still having no independent reason to doubt our own opinion). And so there is some sense in which we have merely possible peers with whom we satisfy the conditions under which EWV does lead to skepticism. If, as Kelly (2005) has argued, the opinions of these merely possible disagreers are just as epistemically significant as the opinions of actual disagreers, then EWV may yet lead to widespread skepticism. In fact, if successful, this challenge will commit EWV to a much broader form of skepticism than the argument in the last section, as for almost *any* proposition we believe, there is some *possible* peer that disagrees with us.

We need to be careful in formulating this new skeptical challenge, however, as there are importantly different versions of it. The simplest version replaces premise (1) of the SCA with:

---

[10] Kelly (2005) presses this worry.

(1′) For each proposition *p* in C, S is aware that it is *possible* for S to have an epistemic peer or superior whose opinion S has no independent reason to doubt with whom S disagrees about *p*.

Barring a radical modal skepticism, it is difficult to dispute that most of us satisfy (1′). Although it is rare that we actually encounter a perfectly symmetrical distribution of opinions on a controversial proposition, such that there really is no independent reason to discount the opinions on one side or the other, we know it is *possible* for this to occur. That is, though epistemic elections rarely result in a tie, a tie is always one of the possible outcomes. This is not a problem for EWV, though. Satisfying (1′) does not entail that we are aware that we are in a disagreement of the sort described in the antecedent of (iii) with any peer or superior (possible or otherwise), but only that we are aware that we *could have* been in such a disagreement. Thus, it does not commit EWV to the claim that our epistemic situation is such that we should split the difference with some peer or yield to some superior, but only that we could have been in such a situation. Since its being possible for us to be in a situation where EWV would require suspension of judgment is no reason to think that we are in such a situation, no skepticism follows.

A stronger challenge is generated by replacing (1) with:

(1″) For each proposition *p* in C, S is aware of a possible epistemic peer or superior whose opinion S has no independent reason to doubt with whom S disagrees about *p*.

This claim is importantly different from (1′) as it requires not merely that S is aware that S could have been in a situation where S had no reason to doubt a peer or superior's contrary opinion, but instead that S is aware that S *is* in such a situation with respect to a possible peer or superior. Nevertheless, instances of (1″) might seem equally difficult to resist; after all, if we want to find

possible peers and superiors who disagree with us about the propositions in C, we need only look to the worlds that make instances of (1′) true—those worlds are at least populated with peers and superiors who disagree with us. Since being aware of an actual disagreer would require us to split the difference, if the opinions of these merely possible disagreers have the same epistemic impact as the opinions of actual disagreers, then if we satisfy (1″), then EWV will require that we split the difference with our merely possible peers and yield to our merely possible superiors. And since we also seem to have merely possible peers and superiors who *significantly* disagree with us, splitting the difference with these peers may in fact lead to skepticism, and not merely slightly reduced confidence, about the propositions in C. Worse yet, merely possible disagreers for *uncontroversial* propositions outside of C, such as the proposition that $1 + 1 = 2$, seem equally available.[11] So, if the existence of these merely possible disagreers commits EWV to skepticism, that skepticism will be quite broad.

But the existence of merely possible disagreeing peers (or superiors) does not commit EWV to skepticism, for the same reasons that the mere existence of *actual* disagreers does not commit EWV to skepticism. That we have some merely possible peer or superior who disagrees with us about p is not sufficient to satisfy (1″); we must also have no independent reason to discount that person's opinion. So, for example, even if the opinions of merely possible peers are given equal weight, awareness of a lone peer in some other possible world who believes that it's not the case that $1 + 1 = 2$ does not make it irrational for us to maintain our belief that $1 + 1 = 2$

---

[11] At least one of us is willing to leave open that there are a few propositions such that there are not even any possible peers or superiors who disagree about them. There may, for example, be propositions about our own mental states such that it is impossible for someone else to be in an equally good epistemic position with respect to them. Such propositions are rare, however.

any more than awareness of a lone *actual* disagreer would. This is because we still have independent reasons to doubt the opinion of this possible disagreeing peer based on what we know about opinions among our other (actual) peers and superiors. As a result, even if we treated this merely possible disagreer as an actual disagreer, (1″) would not commit EWV to any skepticism.

The skeptical challenge can be formulated without this particular liability, however, if we think about entire groups of disagreers, rather than individuals. Consider some proposition *p* for which the balance of opinions of our peers and superiors rationally supports belief to some degree. It seems that there is a possible world at which the balance of opinions of an equally competent and numerous group of epistemic agents rationally supports disbelief, and to the same degree as the former group's opinions supports belief. Since the consensus at this 'peer world' is composed of the opinions of an equal number of epistemic agents in an equally good epistemic position as those in this world, EWV seems committed to granting that collective opinion equal weight, thus requiring us to split the difference and suspend judgment on *p*. Since we can find this sort of peer world for almost any proposition, EWV again seems to lead to broad skepticism. In fact, if we really ought to give equal weight to *all* possible peers, and not just those at one or two worlds, EWV seems committed to the same result. Since for almost any proposition we presumably have infinitely many possible peers (and superiors) who believe it and infinitely many who disbelieve it, the balance of opinions across all possible worlds will be perfectly counterbalanced—for any possible agent believing *p* there will be a possible agent in an equally good epistemic position who disbelieves *p* and vice versa. So again, granting these agents' opinions the weight that they deserve, according to EWV, seems to require suspension of judgment on *p*, thus precluding rational belief or disbelief.

Fortunately for EWV, all versions of this particular skeptical challenge rely on the false claim that being aware of merely possible disagreers has the same epistemic impact as being aware of actual disagreers.[12] Our view is that becoming aware of an actual disagreeing peer or superior gives one higher-order evidence, evidence that at least one person has misevaluated the first-order evidence. When the disagreement is with a peer, neither person is in a better epistemic position than the other, so there is no reason to think that one of them has evaluated the first-order evidence better than the other. This is why both opinions should be given equal weight and why, given no further evidence (such as an independent reason to doubt one of the opinions), both parties should split the difference. When the disagreement is with a superior, that superior is in a better epistemic position, and so there is some reason to think that the superior has not made the mistake. This is why the superior's opinion should be given more weight, and, in the simplest case, the superior's opinions should be adopted. But being aware of merely possible disagreers provides no such evidence. That it is possible for equally competent or even more competent epistemic agents to disagree with us on the same first-order evidence is no reason to think that we have *in fact* made a mistake in evaluating the first-order evidence. Usually, we will agree with our peers and superiors, since we are all reliable judges of evidence; it is only when we actually disagree that we have some reason to think that something has gone wrong—that at least one of us has misread the evidence and our attitudes should be adjusted accordingly. Since it is this feature of awareness of actual disagreers that grants it the epistemic impact that (iii) attributes to

---

[12] For an extended discussion of this point, see Carey (2011).

it, and awareness of possible disagreers does not have this feature, EWV is not committed to our splitting the difference with our merely possible peers.[13]

5. Skepticism and the 'Epistemic Election'

In Section 3, we saw how appealing to an 'epistemic election' can provide a way out of the SCA—the argument that EWV has contingently skeptical consequences about many propositions in politics, religion, ethics, and science. Since the results of 'epistemic elections' regarding such propositions will rarely result in a tie, there will often be some independent reason to discount one of the competitor opinions regarding the disputed proposition. This reason will come from one opinion winning the epistemic election. With such a reason in hand, a non-skeptical attitude toward such propositions will be justified despite their controversial nature. While *which* non-skeptical attitude is justified will depend on the results of the election, appealing to the 'election' provides an escape from the SCA.

---

[13] This position is consistent with claims (i)–(iii) of EWV being necessarily true, but at least one of them will not be true 'across worlds'. Let S1 and S2 be disagreeing peers at different possible worlds w1 and w2. It is consistent with (i)–(iii) being true at both w1 and w2 (and all other worlds) that S1 should not split the difference with S2, since at neither world are the antecedents of (i)–(iii) satisfied—S1 and S2 do not disagree at either w1 or w2. Nevertheless, this situation is incompatible with inter-world versions of (i)–(iii) that allow for subjects from different worlds in their antecedents. A consequence of denying that the opinions of merely possible peers have the same epistemic impact as the opinions of actual peers is that EWV is not committed to these inter-world version of (i)–(iii), either because S2's opinion is not granted equal weight (and so the inter-world version of (i) is false) or because S2's opinion is granted equal weight, but this does not require 'splitting the difference', since S2 does not have a disagreeing opinion at w1 (and so the inter-world version of (iii) is false).

That said, skeptical consequences for EWV re-emerge when the epistemic election is more closely considered. While the results of epistemic elections regarding such propositions will rarely result in a tie, we are often in the dark regarding the results of the epistemic elections for controversial propositions. That is, while the distribution of opinions will typically favor *some* opinion towards *p* over its competitors, it is often the case that we have no reason to think *of any particular opinion* toward *p* that *it* is the 'winning' opinion. Often we will be justified in suspending judgment regarding the results of the epistemic election. But if suspension of judgment is called for regarding the results of the election, then even if there is a 'winning' opinion, we will be unable to utilize this fact to discount some of the competitor opinions. So long as the 'winning' opinion is outside our ken, we will not be able to utilize this critical information to discount competitor opinions and escape from the skeptical consequences.[14] Simply put, without information about which opinion is the 'winner', the skeptical consequences of EWV (contingent as they may be) re-emerge, and unfortunately, we are often without such information.[15]

The argument can be put more formally as follows:

(1) Regarding many of the propositions *p* in C you should suspend judgment regarding the result of the 'epistemic election' regarding *p*.

(2) If (1), then for many of the propositions *p* in C you should suspend judgment regarding *p*.

(3) For many of the propositions *p* in C, you should suspend judgment regarding *p*.

---

[14] This parallels a problem raised by Alvin Goldman (2001).

[15] The skepticism here is more limited since there are some controversial propositions regarding which we are justified in believing of some opinion toward it, that it is the winner of the epistemic election. Yet while this is a more limited skeptical consequence, the skeptical consequence is still quite significant.

Let's begin by examining (1). Premise (1) is made plausible by the fact that you should adopt a skeptical attitude to a couple other propositions:

(a) Regarding many of the propositions p in C, you should suspend judgment regarding the *distribution of opinions* in the epistemic election regarding *p*.

(b) Regarding many of the propositions p in C, you should suspend judgment regarding *how much weight* each of the opinions gets in the epistemic election regarding *p*.

Regarding (a), we are typically simply unaware of *which* opinions are the majority opinions on a great deal of controversial topics. For instance, while we know that it is controversial whether human beings have free will, we are simply unaware of which opinion on the matter is the majority opinion. We know of a number of individuals who have had various competitor attitudes towards this proposition, but our sample is very small and unrepresentative of the whole (both spatially and temporally).

This skepticism is not without its remedy. Various polls and surveys could be taken with the result that we could become justified in believing that certain opinions on these matters are in fact the majority opinions. That said, things are not quite so easy since it seems that it is not only the opinions of our living contemporaries that matter. There have been many great thinkers who were roughly as informed as our contemporaries on the relevant issues. It seems that many individuals who are no longer alive were in as good (or better) of an epistemic position regarding many of the propositions in C. The fact that such individuals are no longer alive does not prevent their attitudes toward these propositions from being higher-order evidence regarding the truth of these propositions. Further, in cases where the relevant first-order evidence remains largely unchanged from their time to ours, their opinions seem to require as much weight as opinions of

our contemporaries.[16] So long as we are justified in suspending judgment regarding the distribution of opinions, we are left with skepticism about the result of the epistemic election, and without a great deal more investigation, that seems to be where we many of us are.

Regarding (b), two other items of ignorance are relevant. Recall that in the epistemic election some votes count for more than others (e.g., expert opinions count for significantly more than the opinions of laypeople). So, even if we knew *how* the opinions were distributed regarding *p*, if we remain ignorant of how to weight the votes, we will remain ignorant of the result of the epistemic election. For many of the propositions in C, we know that some are in a better epistemic position with regard to *p* than others. So, we know that not all votes are to be weighted equally. However, determining *how good* of an epistemic position various people are in with respect to *p* is a very difficult matter. Keeping with our example concerning free will, it may be clear that the votes of metaphysicians are to count for more than the votes of laypeople, but how much more? What about undergraduate philosophy majors? Neuroscientists? How much more weight do those who are doing research on the issue get over those who merely think about it as a hobby?

These are difficult questions to answer. They are answerable. Regarding any one individual we might be able to get a pretty good handle on how good of an epistemic position he or she is in with regard to *p*, but the epistemic election takes *everyone's* epistemic position with respect to *p* into account, since *everyone's* opinion needs a weight. Moving out of ignorance on this matter seems unlikely, even if it is possible.

Another worrisome obstacle which is preventing us from having justified beliefs about the winners of these epistemic elections comes from considerations regarding the independence of

---

[16] See Carey (2011) for more on this point.

opinions.[17] Two agreeing and *independent* opinions about *p* carry more weight than two agreeing opinions about *p* which are not independent (at least other things being equal).[18] Two individuals who both believe that human beings have free will but who went to the same school, took the same classes, had the same advisor, etc. will together give less higher-order evidence for that metaphysical claim than two individuals who don't have such a shared history. Since these two individuals are more likely to be influenced by the same non-epistemic factors, the likelihood of an alternative explanation for their agreed opinion increases, and thus the combined weight of their opinions is less than it would otherwise be. For instance, we are all too aware of the correlations that exist between various philosophical opinions and the graduate school where the individual in question was trained. That independence of opinions will factor into the weight that the combined opinions are to receive is troublesome since we are often unaware of both the degree of independence two opinions have as well as the degree to which that level of independence should affect the weight given to the combined opinions. Independence is indeed a tricky matter. It comes in degrees, and it is unclear precisely how much of an affect it should have. The troubles with getting clear on independence only make it more difficult to avoid having to suspend judgment regarding the winner in the epistemic election.[19]

---

[17] Adam Elga (2007) also raises this issue.

[18] 'Independence' is here being used in a different sense than earlier in the paper. Here it applies to opinions, earlier it applied to reasons.

[19] In addition, disagreements about who the experts are, how much weight various opinions ought to be given, and how independent any two opinions are, will further complicate the matter of obtaining justified beliefs on these matters.

The case for premise (2) is pretty straightforward. Appealing to the epistemic election provided an escape from the SCA. However, if we do not have a justified belief about the 'winning' opinion, then we cannot use the results of the 'election' to avoid the skeptical consequences alleged in the SCA. According to EWV, what we should believe in cases of disagreement is *entirely* a matter of the higher-order evidence—in this case the evidence about the results of the epistemic election. However, if we should suspend judgment regarding the result of the election, then according to EWV we should also suspend judgment regarding the disputed proposition. In other words, if *p* is in C and S should suspend judgment on the results of the epistemic election regarding *p*, then S should suspend judgment on *p* as well. So, as long as we should suspend judgment regarding which opinion is the winner, the mere fact that there is a winner will not save us from these skeptical consequences. While the existence of a winner might mean that there is some reason out there to be had which would allow us to discount the competitor opinions, unless we have some sort of access to that reason, it won't allow a route of escape from the SCA.

One might think that having a justified belief about what 'winning' opinions is, is easier to come by than we have been making it out to be. Can't one generalize from one's own opinion about *p* to the winning opinion about *p*? Suppose that Smith takes himself to be a rational individual, and knows that he believes *p*. Can't he generalize that the attitude of *belief* is the winning opinion toward *p* in the epistemic election since most people are rational, and so most people would adopt the attitude he has toward *p*? If Smith thinks about the claim that the Earth is further from the Sun than Mars, and sees that he believes it, can't he also rationally conclude that belief is the winning opinion toward this proposition in the epistemic election? After all, isn't it rational for him to conclude that anyone in his epistemic position will believe this, and anyone who is in an even better epistemic position will believe it as well?

While reasoning in such a way might be legitimate regarding many propositions, it often won't be regarding the propositions in C, at least not once the individual in question is aware of the controversy surrounding these propositions. If one knew nothing of the controversy regarding a proposition in C, then such a generalization may well be warranted, yet for better or worse, informed individuals will be aware of the controversy surrounding these propositions. The awareness of the controversy defeats such an inductive conclusion. Once one is aware of the controversy surrounding such a proposition, it is no longer reasonable to suppose that someone in the same epistemic position would hold the same opinion, or that those in better epistemic positions would hold it either. Being aware of the controversy is to be aware that there are significantly many others in equally good and better epistemic positions who hold a different opinion toward the proposition in question. So, while such an inference may present another way in which EWV can avoid these skeptical consequences (it will not have skeptical consequences for those who are unaware of the controversy surrounding the propositions in C), it fails to avoid the skeptical consequences for those of us who are aware of the controversy surrounding these propositions.

In Section 3, we saw that appealing to an epistemic election could give us a reason (and one independent from the disagreement itself) to privilege one opinion on the disputed matter. But, to have such a reason, we must be justified in believing that one opinion on the matter is the winner of the epistemic election. What we have seen is that (at least right now) we don't have such a reason, at least regarding many of the propositions in C. While there will typically be a victorious opinion in these epistemic elections, the fact that we should suspend judgment regarding what that opinion is prevents us from having such a reason to discount any of the conflicting opinions.

Further, appealing to the epistemic election yields yet another argument that EWV has skeptical consequences in the contingent circumstances of controversy that we find ourselves in. Recall that according to EWV what one should believe in a case of disagreement is entirely a matter of the higher-order evidence. What we have seen is that regarding many of the propositions in C we should suspend judgment as to what the higher-order evidence supports— we should suspend judgment as to the result of the relevant epistemic election. However, if what we should believe about disputed proposition *p* is entirely a matter of the higher-order evidence, and we should suspend judgment regarding what that evidence supports, then we should also suspend judgment regarding *p* as well. While appealing to epistemic elections does allow EWV to avoid skeptical consequences for agents who have good reasons to believe in particular outcomes for those elections, we do not seem to have such reasons.

6. Skepticism and Nearby Disagreements

Even in cases where we do have sufficient evidence to justify belief in some particular outcome to the epistemic election, considerations about disagreement in *nearby* possible worlds may still require suspension of judgment. We saw in Section 4 that no amount of evidence for the mere possibility of disagreement is sufficient to generate the sort of higher-order evidence that makes suspension of judgment plausible in cases of actual disagreement. There are, however, other ways in which awareness of the opinions held by peers and superiors in other possible worlds *can* give us higher-order evidence. Though the mere fact that there is a peer world where opinions are significantly different from the opinions we have good reason to believe obtain in the actual world does not yield the requisite sort of higher-order evidence, we can have good reason to believe that there are many such peer worlds that are in some important sense 'nearby'. If so, then

we would have some important evidence about the higher-order evidence we have concerning the relationship between the first-order evidence and the target proposition coming from the distribution of opinions in the actual world.

Consider, for example, the dispute between compatibilists and incompatibilists about free will and determinism. Both positions have at least some intuitive support, and there are valid arguments in favor of each with no obviously false premises. There are exceptional philosophers defending each position, but for reasons of the sort described in Section 3, the weighted overall balance of opinions is probably not perfectly counterbalanced. Supposing we could overcome the problems in the last section enough to reasonably believe in a particular result to the actual epistemic election, unless we have independent reasons to doubt some of the opinions, EWV would therefore require that we adjust our attitude in the direction of whichever position won the election. Considering further facts about the nature of this disagreement, though, might lead one to reasonably question whether the results of the epistemic election accurately reflect what the first-order evidence supports. For one thing, the known evidence that bears on the truth of compatibilism is extensive, such that it is reasonable to think that even an expert could easily have trouble keeping track of every relevant item of evidence. And there seems to be no position that fits well with *all* of this evidence; even the staunchest compatibilist will have to admit that there are some considerations that clearly seem to favor incompatibilism, and vice versa. As a result, correctly evaluating the relevant evidence requires correctly weighting the epistemic importance of conflicting items of evidence, which even experts on the subject may not be adept at.

Furthermore, there are a significant number of experts who have concluded that this evidence supports compatibilism and a significant number who have concluded that it supports incompatibilism. This suggests either that the evidence is difficult to evaluate, or that even those

best at evaluating it are not especially good at evaluating it, since a significant number of them have apparently misjudged it. Depending on exactly what evidence one has about the free will debate, these types of considerations might make it reasonable to believe not only that there is some possible world where the epistemic election has significantly different results, but that there are many such worlds intuitively nearby to the actual world. And we can imagine having further evidence that would make it even more reasonable to question the epistemic significance of the actual balance of opinions.

Suppose, for example, that the known history of the free will debate was one in which the weighted majority opinion regularly shifted between compatibilism and incompatibilism without any significant changes in the arguments, examples, or other shared first-order evidence. Having noticed the weighted balance of opinions changing while the evidence being evaluated remained the same would give a good reason to think that the opinions could easily have been different, even with the same first-order evidence.

For disagreements about controversial propositions of this sort, where it is reasonable to believe that there are many nearby worlds where the balance of opinions is significantly different even though the first-order evidence is not, yet another skeptical challenge arises. Consider some body of first-order evidence E and some proposition p such that the actual consensus opinion is that E supports p. Assuming that evidential relations are necessary,[20] either E supports p at all possible worlds or E does not support p at all possible worlds. So, if there are many nearby worlds where the consensus is that E does not support p, then either the consensus opinion is

---

[20] We find this assumption plausible, though a weaker one will also do. So long as evidential relations remain the same at all nearby worlds, the same problem can be generated, even if there are distant worlds where the evidential relations are different.

actually incorrect or there are many nearby worlds where it is incorrect. But this seems to undercut, to at least some degree, the reason for believing that E supports *p* based on the consensus opinion, since that opinion either is mistaken or easily could have been mistaken. Put another way, if there is good reason to believe that there are many nearby worlds where whether E supports *p* remains the same, but the consensus opinion about whether E supports *p* is different, then the consensus opinion is no longer a good indicator of whether E supports *p*. If we think of the epistemic election as a sort of epistemic measuring device, then in this kind of situation, it is a device whose readings do not track what the evidence supports, and so the results of the election are not a good reason to think that E in fact supports *p*. In such situations this higher-order evidence from the consensus opinion is undercut.

Given what we have said about higher-order evidence, this result is not surprising. In a simple case of peer disagreement, the subject's awareness of a disagreeing peer whose opinion there is no good reason to doubt undercuts the subject's reason for believing that the first-order evidence supports the target proposition. The effect here is similar, but at a higher level of evidence. *Prima facie*, learning that the weighted balance of opinions of one's epistemic peers and superiors is that E supports *p* provides some evidence that E supports *p*. But learning that this balance of opinions could easily have been otherwise, that there are many nearby worlds where the same epistemic agents evaluate the same first-order evidence and reach a significantly different consensus, undercuts the support that evidence about the results of the epistemic election would ordinarily confer on the proposition that E supports *p*. So, evidence that opinions could easily have been different is *third-order* evidence: evidence about one's (second-order) evidence about whether the first-order evidence supports *p*. It is not evidence that bears directly on *p*, nor even evidence that bears directly on whether E supports *p*, but rather evidence that bears on whether the evidence from the epistemic election supports that E supports *p*. In fact, evidence

26

of many nearby worlds with a different consensus is just one type of third-order evidence that might undercut the epistemic weight of our evidence about the actual consensus. We might, for example, have reason to believe that many peer and superior opinions were not formed based on the evidence, but because of some other theoretical commitment, as a result of wishful thinking, or on a whim. Or we may have evidence that only those already committed to believing *p* are interested enough to become experts regarding *p*, and so there are no experts who disbelieve *p*, but not because the evidence clearly favors *p*. Evidence of any of these issues in a sufficient number of those polled in the epistemic election would be a good reason to think that the results of that election do not accurately track what the first-order evidence supports.

The route to skepticism is familiar, but with an extra step. If S has sufficiently good reason to think that the balance of opinions could easily have been different, then S ought to suspend judgment on whether the balance of opinions supports that E supports *p*. But, since this evidence swamps all other evidence S has about whether E supports *p*, S should then suspend judgment on whether E supports *p*. As a result, S should suspend judgment on *p*, since it is not reasonable for S to believe that the evidence bearing on *p* (that is, the first-order evidence) supports *p*, nor is it reasonable for S to believe that that evidence supports ~*p*.

But there is little reason to think that skepticism of this sort is as widespread as the skepticism that threatened in Section 4. It does not seem to arise, for example, for non-controversial propositions. There are those who believe that the Sun revolves around the Earth, but we have no reason to think that the balance of opinions about this proposition could easily have been different with no change in the first-order evidence. The stability of expert opinion over time that the Sun does not revolve around the Earth and the lack of any significant number of experts who disagree both give good reason to think that in our 'modal neighborhood' there are at most very few worlds where the weighted balance of opinions is significantly different.

There are, of course, many possible worlds in which the consensus is that the Sun does revolve around the Earth, even though the same first-order evidence is possessed, but these worlds are not nearby in the relevant sense—they are worlds where, for example, the experts are less competent than in this world, or they use methods different from those our experts actually use, or they form their opinions on the basis of something other than the evidence.[21] The existence of such worlds does not seem to undermine our evidence in the same way that evidence of many nearby worlds where an equally competent epistemic community draws a different conclusion from the same body of evidence does. Nor does this type of skepticism threaten for all controversial propositions; we may have good reason to think that there is significant disagreement on some proposition without thereby having a reason to think that the balance of opinions could easily have been different. For example, it is, at least in some sense, controversial whether childhood immunizations cause autism. If all we knew was that the results of the epistemic election favored the belief that immunizations do not cause autism, then it would be reasonable to believe that they do not, even though we know there is some disagreement. Suspension of judgment would only be required in the face of further evidence about the consensus of our peers and superiors— that they could easily have come to the opposite conclusion on the basis of the same evidence. Since we have no such evidence, there is no threat of skepticism.

7. Conclusion

We have seen that the standard arguments that EWV is committed to skepticism about controversial propositions based solely on our knowledge of actual or possible disagreeing peers and superiors are unsound. Nevertheless, our lack of information about the distribution of

---

[21] Hilary Kornblith (2010) makes a similar point about Gödel's proof of the incompleteness of arithmetic.

opinions on controversial propositions, as well as how to weigh them, and our reasons to doubt that this distribution of opinions is a good indicator of what the evidence supports lead to skepticism about at least many of the same propositions. This skepticism is less problematic for defenders of EWV, however, even if its scope is nearly as broad. Although we should suspend judgment more often than we do, this is not simply a result of knowing that there are those who do or could disagree with us; it is a result of our having insufficient evidence to justify belief in what our first-order evidence actually supports.

References

Carey, B. (2011). "Possible Disagreements and Defeat," *Philosophical Studies* 155: 371-381.

Christensen, D. (2007). "Epistemology of Disagreement: The Good News," *Philosophical Review* 116: 187–217.

_____. (2009). "Disagreement as Evidence: the Epistemology of Controversy." *Philosophy Compass* 4(5): 756-767.

Elga, A. (2007). "Reflection and Disagreement" *Noûs* 41(3): 478-502.

Feldman, R. (2006). "Epistemological Puzzles about Disagreement," In Steve Hetherington (ed.), *Epistemic Futures*. Oxford University Press.

_____. (2007). "Reasonable Religious Disagreements," In L. Antony, (ed.), *Philosophers without Gods: Meditations on Atheism and the Secular Life*. Oxford University Press.

Feldman, R., & Warfield, T. (eds.). (2010). *Disagreement*. New York: Oxford University Press.

Goldman, A. (2001). "Experts: Which Ones Should You Trust?" *Philosophy and Phenomenological Research* 63(1): 85-110.

Kelly, T. (2005). "The Epistemic Significance of Disagreement," *Oxford Studies in Epistemology* 1: 167–96.

———. (2010). "Peer Disagreement and Higher-Order Evidence," 111–74 in Feldman & Warfield (2010).

Kornblith, H. (2010). "Belief in the Face of Controversy," in Feldman & Warfield (2010).

Lackey, J. (2010). "What Should We Do When We Disagree?" *Oxford Studies in Epistemology* 3: 274–93.

Matheson, J. (2009). "Conciliatory Views of Disagreement and Higher-Order Evidence," *Episteme* 6: 269–79.

———. (manuscript). "Disagreement: Idealized and Everyday."