PRACTICING CONCILIATION: TOWARDS A PRACTICAL APPLICATION OF THE EQUAL WEIGHT VIEW

A thesis submitted in partial fulfillment of the requirements for the degree of

Master of Arts

of

Rhodes University

by Danyel Hartley

March 2021

Abstract

The Equal Weight View is a frequently discussed position in the philosophy of disagreement. It holds that when someone disagrees with an epistemic peer, they should adjust their belief to be closer to their peer's belief. While the reasons for adopting this response to disagreement have been debated, there has been less discussion about its utility as a tool for handling real-world disagreements. In this thesis I pursue a version of the Equal Weight View which is useful in practice. I argue that traditional applications of the Equal Weight View do not reflect its underlying principles when they are used to resolve real-world disagreements. I develop an idealized application of the Equal Weight View that addresses the problems traditional applications face in real-world scenarios. Unfortunately, addressing these problems results in an application that is unrealistically cognitively demanding. The application trades being insensitive to the environment it would be used in for being insensitive to the limits of its user. I suggest that we might be able to save the idealized application and work around those limits by either externalizing or simplifying the most demanding aspects of the application. Externalization is best achieved by making use of some kind of computer assistance. Simplification involves replacing taxing data tracking and computation with heuristic methods. While neither approach is ideal, I argue that both get us closer to resolving disagreements in accordance with the underlying principles of the Equal Weight View than traditional applications do.

Table of Contents

List of Figures:

Acknowledgments

I want to thank my supervisors, Therese Dewhurst and Richard Flockemann, for their incredible patience, insight, and guidance over the course of writing this thesis. Most of all I would like to thank them for their encouragement, without which this project might never have been finished. I also want to thank Ward Jones, Francis Williamson, and Marius Vermaak for their feedback and discussion on earlier versions of these ideas.

Introduction

What am I to do when I disagree with my epistemic peer? My 'epistemic peer' being someone who is just as clever, informed, or likely to be right as I am. Does knowing about such a disagreement affect what I should believe? This is the question at the heart of the problem of peer disagreement (PPD).

The Equal Weight View (EWV) is a solution to the PPD defended by Richard Feldman, David Christensen, and Adam Elga.[1] In essence, the EWV holds that when I learn that my epistemic peer disagrees with me, I should adopt a conciliatory position and shift my belief towards my epistemic peer's. I ought to reconcile our positions by assigning equal weight to both of our views. To 'assign equal weight' means to think of each belief as being equally likely to be correct or mistaken. If my epistemic peer believes that God exists but I believe that God does not exist, assigning equal weight to both beliefs would mean adopting agnosticism (i.e., believing that it is as likely that God exists as it is that God does not exist).

With this thesis, my objective is not to argue in favor of the EWV as a solution to the PPD. I am going to assume that the principles underpinning the EWV are correct and that it is a theory we are interested in using. Specifically, I want to use the EWV to address real-world cases of disagreement. The EWV has been developed as a solution to the problem of peer disagreement. But I think that the EWV also has clear implications for how we should resolve disagreements with people who are not our epistemic peers. In fact, the PPD is likely to be a fairly rare problem. By the definition of the PPD I will be using it is not as rare as it could be,[2] but it is still uncommon. Yet, if we think that the EWV tells us how to respond to peer disagreement, surely it also suggests how to resolve disagreements with epistemic superiors and inferiors. Thus, while the PPD might be rare, how we respond to it has implications for *every* disagreement we could find ourselves in. The utility of the EWV, when it is restricted to being only a solution to the PPD, is needlessly limited. I will discuss the PPD in some detail, but ultimately, I intend to use it as a springboard for an application of

[1] See (Christensen, 2007; Elga, 2007; Feldman, 2007) among others.

[2] Indeed, by Thomas Kelly's (2005) definition it is almost impossible to meet an epistemic peer, as I shall discuss in section 1.4.3.

the EWV that can address disagreements with not just our epistemic peers but all people, something that standard applications of the EWV fail to do.

The EWV is usually presented and discussed as a solution to an abstract thought experiment. These experiments usually involve isolated cases of peer disagreement. However, a key property of disagreements in the real world is that they aren't isolated. Focusing on the EWV as a solution to only isolated instances of disagreement, rather than as a cognitive tool that will be used by normal people to resolve disagreements over the course of their lives, has resulted in an unfeasible position. I argue that the EWV, as usually presented, is not a particularly *useful* solution to cases of disagreement outside of unusually isolated ones. Usually, the prescribed application will fail to process evidence without violating the EWV's own underlying principles when we attempt to use it outside of isolated scenarios. I argue that these applications cannot be how we prescribe following the EWV.

Thus, my goal is to develop a version of the EWV that is able to both reflect the principles the theory is committed to and be practically useful in day-to-day life. I call this the Practical Equal Weight View (PEWV). Developing an application of the EWV that is both aligned with its principles and genuinely useful turns out to be a difficult task. To achieve both goals will require some compromises.

In Chapter 1 I clarify the goals of the thesis and map out the foundational concepts I will be working with. I discuss what the objective of finding a *practically useful* application of the EWV means. I am looking for a solution that works with disagreements as they appear in the real world, rather than one that works well in unrealistic and unlikely theoretical cases. I argue that a practical solution needs to be sensitive both to the internal properties and limits of the person using it, as well as the external real-world conditions it is meant to be used in. The EWV is usually intended as a solution to the problem of peer disagreement, so I spend some time outlining just what this problem is.

In Chapter 2 I examine the standard EWV as a solution to the PPD as it would occur in the real world. I discuss the principles underlying the view, and cases where it appears to work both well and poorly. In this chapter we find that the standard EWV simply does not *work* when applied to the type of real-world disagreements we find outside of isolated thought

experiments. I propose a view that reflects the principles of the EWV in real-world cases of disagreement and call it the Idealized Equal Weight View (IEWV). Though I only develop the IEWV in Chapter 3, I analyze major problems that have been identified with real-world use of the standard EWV that the IEWV needs to address here. The EWV is surrounded by several competing positions with their own merits. Because my aim is not to argue that the EWV is a better solution to the PPD than its competition, I will not discuss these views in detail. An exception is the Steadfast View (SV) which holds that learning that an epistemic peer disagrees with me has no bearing on the rationality of my beliefs. Thus, I can and should continue believing as I did prior to the disagreement. I compare how the EWV and SV handle scenarios from the literature to get at what is intuitively attractive about both views. I aim to incorporate some of the apparent strengths of the SV into the IEWV.

In Chapter 3 I develop an application of the EWV that avoids the problems discussed in Chapter 2 – the Idealized Equal Weight View (IEWV). By changing the way beliefs are processed, stipulating which beliefs should be used as inputs, and creating a ledger to record those beliefs in, the IEWV achieves part of my goal. That is, it both works and accurately reflects the principles of the EWV in real-world cases. As a result, the IEWV can be used in any instance of disagreement, not only with the PPD. However, the IEWV is an *idealized* view, not a *practical* one. It can be used to resolve real-world disagreements, but it isn't something real people would want to use. While it is possible for a human to use it, it is too cognitively demanding for us to really incorporate it into our day-to-day thinking. So, the IEWV would function correctly in a real-world context (unlike standard EWV applications) but is not adequately sensitive to those internal human limitations.

In Chapter 4 I use IEWV as a model for developing a Practical Equal Weight View (PEWV). I use the IEWV because, unlike standard applications of the EWV, it actually reflects the principles of the EWV when resolving real-world disagreements. I argue for two possible paths to the PEWV. To reduce the cognitive cost of using the view, the processes developed for the IEWV can be *externalized* or *simplified*. A path towards externalization would be to develop a computer program that handles the parts of the IEWV which are too demanding. There are many advantages to this approach, but it does commit us to rely on and interface with an external system. The second option, simplification, would involve utilizing heuristics to replace overly demanding functions with more efficient ones. Neither solution is ideal,

however, both bring us closer to applying the principles of the EWV in the real world than the standard EWV does.

***

Chapter 1. The Problem of Peer Disagreement

1.1 Introduction

Does the fact that my epistemic peer disagrees with me mean I should change what I believe? This question is at the heart of the problem of peer disagreement (PPD). According to the Equal Weight View (EWV), being aware of a disagreement with an epistemic peer is a very good reason to revise my beliefs. On the other hand, the Steadfast View (SV) holds that the mere fact of disagreement itself shouldn't affect what I believe at all.

In this thesis, I assume that the Equal Weight View (EWV) is the correct solution to the PPD. Or rather, I assume that the principles the EWV is based on are sound. I think that standard methods of applying the EWV can't avoid violating these principles outside of isolated abstract cases of disagreement. I will refer to the various attempts to use the principles of the EWV to resolve disagreements as "applications". The version of the EWV most commonly discussed in the literature is what I call the 'standard application' or 'standard EWV'.

My primary goal with this thesis is to arrive at an application of the EWV that could be applied in the real world without violating its own principles. I intend to achieve this goal by developing a practical application of the EWV that real human beings could make use of. I call this application the Practical Equal Weight View (PEWV). All of the discussion that follows is in this thesis is in the service of developing the PEWV.

There are serious problems with the standard EWV that the PEWV needs to address. But before discussing where the EWV goes wrong, or how it could be repaired, I need to establish exactly what the problem it needs to solve is. Thus, getting to grips with just what the PPD means is my focus in this chapter. In section 1.2, I begin to outline the PPD with an example case. In section 1.3, to understand what the parameters for a practically useful solution to the PPD are, I discuss what it means for a solution to be practically useful in the real world. Finally, in section 1.4 I provide and discuss the definitions I will be using to make sense of the PPD.

1.2 The Problem

Let's consider the case of a young university student, Hob. Hob is interested in the topic of climate change. He has a basic understanding of how greenhouse gasses are supposed to cause global warming, which he acquired during his school science classes. More recently, Hob's university education has been in the humanities, and he has ignored science in his personal studies. Thus, Hob, like most people, lacks a sophisticated understanding of contemporary climate science. Hob decides that he should do some research on the internet, focusing on the question of whether anthropogenic climate change is real. That is, is climate change caused by humans? Over the course of his research, he discovers a series of news articles where he reads that most scientists agree that anthropogenic climate change is a real problem. These articles seem to Hob to be credible. He attempts to read some of the scientific studies that have been published but quickly gives up on understanding them - this complex work is too far outside his areas of expertise. Finally, he reads some articles written by climate change deniers. Hob finds the deniers' arguments to be unconvincing. At the end of his research Hob weighs up all of the evidence he collected and is close to certain that climate change is caused by humans.

Later that evening, Hob meets up with a friend, Fae. While Hob was doing his research, he sent all the articles he read to Fae, planning to discuss them later. It is important to note that Hob considers Fae to be fairly smart - just as smart as Hob is, in any case. Hob respects Fae as a thinker and takes her beliefs seriously. I will say that, as far as Hob is concerned, Fae is his *epistemic peer*, his intellectual equal. However, when they begin to discuss Hob's research into climate change, Hob learns that Fae thinks that the articles they both read show that it is very *unlikely* that climate change is caused by humans.

So, given that Hob finds himself disagreeing with Fae, whom he considers an epistemic peer, how should he respond? There are many courses of action available to us when faced with disagreement, but for the purposes of this thesis I am not interested in considering all of them. Some of these actions might change the circumstances of the disagreement. For example, Hob might enquire into Fae's background knowledge on the subject (her total pool

of evidence) which might change how Hob views their epistemic relationship.[3] Perhaps Fae has read something that Hob has not, or vice versa. After exchanging this missing information, they might find that they agree that climate change is or isn't caused by humans. Alternatively, Hob might pursue an argument with Fae, hoping to resolve the disagreement by convincing her to change her mind (or by having his mind changed in the course of their argument). I don't doubt that often these actions are appropriate. Indeed, for epistemically honest thinkers a good faith discussion aimed at working out why they disagree could be the very best thing to do. However, such actions are not always available to us. Sometimes we don't have the time or energy to have these discussions or debates. Furthermore, when we do have these discussions, they often fail to bring us any closer to an agreement or even further entrench our disagreement.

I think we can all agree that disagreements can be frustratingly difficult to resolve. Indeed, those disagreements found among philosophers seem to be of an especially stubborn kind.[4] The ubiquity of disagreement in our world makes it clear that we have not yet discovered a method for reliably resolving disagreements. If someone had discovered this method, we would all agree with her. Thus, as most of them aren't going to disappear, it is worth asking if the mere presence of a disagreement is itself a threat to what we believe. Or, in a more positive light, can the fact of a disagreement point us towards better beliefs? Should Hob maintain his belief as it was before learning of his disagreement with Fae, or should he revise his belief in some way? Does learning that he disagrees with Fae mean anything for Hob's belief that climate change is caused by humans? I will state the problem of peer disagreement in Hob's case as follows:

> Hob possesses the belief that climate change is caused by humans based on a body of evidence (ranging from his high school science education to his recent research). How should Hob's belief that climate change is caused by humans be affected when he learns that Fae, whom Hob considers an epistemic peer, disagrees with him?

---

[3] For example, Hob might come to think that Fae is in fact not as informed as he is on this topic, and so no longer think of her as his epistemic peer. However, as I will discuss in section 1.4.3, I don't think that sharing the same pool (or even equivalent pools) of evidence is necessary for establishing epistemic peer-hood.

[4] (Sosa, 2013)

Hob's case can be used as the base to derive this abstract formal statement of the problem of peer disagreement (derived from Brian Frances' description)[5] which I will be using in this thesis:

> *A* possesses some belief *P* based on evidence *E*. How should *P* be affected when *A* learns that *B*, who *A* considers an epistemic peer, disagrees about *P*?

1.3 A Practical Solution

In philosophy there is a risk of becoming preoccupied with what is true or correct only in an abstract theoretical domain. The resulting arguments are certainly valuable and interesting, but they can be divorced from practical questions about how we should think and act. Thus, we find ourselves left with unwieldy theoretical solutions to what are actually very real problems. If we want to solve problems in real life, we need *practical* solutions. To see just how theory and practice can come apart let us consider how to help Hob with another problem – running a marathon.

We aren't going to be Hob's personal running coaches; we're just going to suggest a training plan for him to follow. So, what is the best marathon training plan for Hob? There is a lot to consider when trying to answer this question. What properties would make one plan better than another? To some extent that will depend on Hob's goal. To keep things simple, let's say that his goal is to complete a marathon quickly. A big part of what would make a plan practically useful is that it actually gets Hob to achieve this goal.

Over the years many marathon training plans have been developed. These plans vary in several ways. To mention just some of these variables: some plans can take longer to complete, utilize different exercises, specify different running forms, place greater emphasis on diet, involve more cross-training, or simply be more intense than others.

It is tempting to think that by looking at how contemporary elite marathoners train we could avoid getting distracted by these variables. Surely Eliud Kipchoge, an Olympic gold medalist in the marathon, follows one of the best training plans around? Indeed, given that Kipchoge

[5] See Frances (2014, p. 83).

holds the record in the marathon we could say that his training plan is *the best plan*. Following this plan made Kipchoge complete a marathon faster than anyone else ever has. So, given Hob's goal, it has produced better results than any other plan. But there are many ways in which Kipchoge is unlike Hob (or the overwhelming majority of people, for that matter), and there is no doubt that some of those differences enable him to train in a way that Hob could not. Kipchoge's training plan demands a great deal of time, commitment, and (most importantly) physical strength and endurance. But Hob is new to running, works full time, and has a bad knee. In fact, Kipchoge's easiest run on a rest day is both longer and faster than Hob has ever run in his life. Thus, it is almost certain that if Hob were to try to train like Kipchoge, he would end up abandoning the plan or injuring himself long before race day. So, While Kipchoge's plan is an excellent plan for an elite marathoner, it would likely be worse than training with no plan at all for Hob[6] - without a plan, Hob is at least less likely to injure himself. The problem is that Kipchoge's training plan, despite being a good plan for those who can meet its demands, is simply too demanding for most people, and is definitely too demanding for Hob.

Analogous solutions in philosophy are ones that are theoretically sound and yet demand too much from us to ever be properly put into practice. Utilitarianism is a classic example, sometimes requiring calculations that are too complicated to be completed before we need to decide how to act – as is acknowledged by utilitarians.[7] For a utilitarian determining how to act means performing the hedonic calculus, wherein we work out how much pleasure will come from an action – this can be taxing work. While I might find utilitarianism to be a compelling moral theory, it can't always be used to resolve moral quandaries in the real world – sometimes it simply takes too long to do the necessary calculus. To make utilitarianism useful in such circumstances the application of the theory needs to be faster. If the application becomes faster, it can be used to solve more moral problems as they arise in the real world and is thus more practical. Mill argued that we could do this by using the greatest happiness principle to determine our laws.[8] In our day-to-day lives we follow those laws, which, unlike the greatest happiness principle, provide quick and clear guidance. I

[6] See Adam Morton (2012, p. 7) on "The approximation fallacy".

[7] (Mill, 1863, p. 23)

[8] (Mill, 1863), see also (Bentham, 1983).

intend to show that the EWV is a lot like utilitarianism in this regard. My interest in the philosophy of disagreement is rooted in that it might provide a clear course of action for how we should think and behave in our ordinary human lives. We all encounter disagreements, and I think it is important that we can respond to those disagreements in the best way possible. Later, in section 2.5.5, I argue that the EWV has implications for how we respond to all disagreements. If it is to be a practical solution to the problem of peer disagreement and to disagreement in general, the EWV can't just be theoretically sound; it also needs to be something that we can actually *use* when we encounter those disagreements.

Hob should clearly avoid Kipchoge's training plan - he simply isn't physically capable of meeting its demands. Beyond ignoring our basic abilities, there are other ways that plans can be poorly suited to real-world use. Imagine a plan that has been generated for Hob by a computer. This plan is incredibly detailed, laying out how every *second* of Hob's day should be spent. The computer is sensitive to Hob's physical limits, so it will always prescribe activities that he is capable of safely performing. We could even say that if it were followed perfectly this plan would make Hob as fast as he could possibly be by race day. This sounds like a pretty good plan, but most of the day is allotted to doing incredibly boring activities. Worse, the program constantly updates Hob on exactly what movements he should be making. In fact, it is unlikely that Hob (or most people) could follow this plan very well given how tedious it is – he would get too bored and decide to do something else instead.[9] There are limits to what people are willing to put up with to achieve their goals. Hob only cares so much about this marathon, and so even though he could use the computer-generated plan, to him it isn't worth the effort and boredom. Hob would quickly stop following the plan perfectly or abandon it all together. Perhaps this plan is also only effective when it is followed perfectly – as abstinence is the best contraceptive strategy. We might worry that such a plan, while a sound solution in theory, would in practice rarely be enacted perfectly. A more forgiving, if slightly less effective, plan might, on average, produce better outcomes because it would be more closely followed.

In this thesis I am looking for a version of the EWV that accounts for the practical dimensions of human life so that it can be of use to humans. Solutions to the PPD quickly

[9] See Morton (2012, p. 24) on "Prescriptions vs programmes".

stop being practically viable when they ask more of us than we are capable of or willing to give. The plans we have looked at for Hob are too strenuous for him, or expect him to have a higher level of fitness going into training than he does, or are too monotonous to keep him engaged. At this point it should be clear that part of what makes something a good *practical* solution is that it is well-tailored for its user. That means being sensitive to all of the internal features of the user, from their capacities to their desires. Hob only cares so much about running a marathon quickly, and in real life we only care so much about what we believe.[10] We want the best results for the least effort, and there is a limit to how much effort we are willing to expend. The PEWV must achieve its goal as best it can without asking so much of us that we're unlikely to use the solution.

The practical problems I have discussed so far arise because the solution doesn't 'fit' its user. I am going to call them 'internal practical problems' because in a sense they come from *inside* the user. That is, the solution isn't sensitive to the real-world properties of its user. I want to contrast 'internal practical problems' with 'external practical problems'. External practical problems arise where the real-world nature of *the problem itself*, or the environment or context that the problem arises in, isn't properly addressed. Regarding our attempt to find a training plan for Hob, this might mean that the solution doesn't account for things like the altitude or weather Hob is going to be training in, or even the details of the race he actually wants to run. External practical problems occur when there is a mismatch between the idealized conditions in our heads and the real-world conditions he is actually going to be running in. Regardless of how capable Hob is, the training plan we recommend could fail him if it doesn't account for him training at high altitude, in very hot weather, and for a very hilly race. Indeed, a plan that fails to account for these variables might not make Hob perform any better in his race, and thus completely fail to help Hob achieve his goal. The PEWV needs to be a solution that is, so to speak, both internally and externally sensitive.

[10] For example, while I care quite a bit about my broad political theory being correct, I don't care so much about this belief to have built my life around inquiry into it. A few hours here and there, but that's about as far as it goes. I care much less about most of my other beliefs. Of course, I would like for them to all be correct, but I'm willing to accept a higher chance of being wrong if it means I don't need to spend all day researching the mundane.

I think that the EWV runs into many external practical problems because the idealized cases it is designed to solve are very different from the cases of disagreement we encounter in the real world. As a result, the standard EWV does not *work* in the real world. I discuss these problems in detail in section 2.6. In Chapter 3 I develop the Idealized Equal Weight View (IEWV) to overcome these problems, but unfortunately this view is fraught with *internal* practical problems. For the IEWV to work in the real world, we need to find ways to overcome its internal practical problems, and in doing so develop the PEWV. I attempt to do this in Chapter 4.

1.4 Definitions

Consider that abstract statement of the PPD again:

A possesses some belief *P* based on evidence *E*. How should P be affected when *A* learns that *B*, who *A* considers an epistemic peer, disagrees about *P*?

In order to develop a practical application of the EWV it needs to be clear just what the EWV is trying to achieve and why it has those objectives. The EWV was developed as a solution to the PPD, so understanding it requires coming to grips with just what the PPD is about. This will require defining the concepts involved in the PPD. This part of the thesis arrives unusually late because I must explain my definitions by way of reference to my practical concerns. I discuss just what is meant by 'belief', 'evidence', and 'epistemic peer' below. There is debate about how some of these concepts ought to be defined in the literature. In these cases, I will offer some arguments for why I favor one definition over another (bearing in mind my goal of arriving at the PEWV).

1.4.1 Belief

First, what is a 'belief'? I am going to take a belief to be an attitude about the truth value of a proposition. So, believing that it is sunny outside just means taking the proposition [it is sunny outside] to be true, rather than false. Hob and Fae *disagree* because they have different

attitudes towards the truth value of the same proposition (Hob thinks that [climate change is caused by humans] is true, while Fae thinks it is false).[11]

I assume that the range of attitudes available to us about the truth value of a proposition is more granular than merely taking it to be true or false. We can also be agnostic about a proposition; by which I mean that we think it is as likely to be true as false.[12] Furthermore, there can be subtle differences in our attitude about how *likely* a proposition is to be true or false. That is, we can be more or less certain that a proposition is true or false. I might be quite sure that it is daytime but be less certain that it's sunny. So, there is a sense in which we believe to different degrees by thinking that *P* is more or less likely to be true. I will use the standard notation of measuring these granular attitudes, called 'credences', between 0 and 1. A credence of 0 means being certain that a proposition is false, 1 means being certain that it is true, and 0.5 means being agnostic - i.e., thinking it is as likely to be true as false.[13] By this model, Hob and Fae would disagree if their respective credences in a proposition were at all different. Even if that difference was *very* small, say Hob had a credence of 0.45 and Fae one of 0.44, they would still disagree about what the appropriate credence that *P* is.[14] Thus, I take it that to 'disagree' is just to have different credences about the truth value of the same proposition.

1.4.2 Evidence

By 'evidence' I mean something taken to bear on a proposition's truth value. I assume an evidentialist position, so only my own mental content can be evidence for *P*.[15] The structure and role of evidence are very important in the EWV in general, and the applications of the view I develop here in particular. I will return to the topic in section 2.4.

---

[11] (Matheson, 2015, p. 7)

[12] This position is sometimes described as 'withholding judgment' about *P*, which, depending how you conceptualize believing, may be different to judging *P* to be as likely to be true as false. Whatever the case, I think that both attitudes lead to very similar behaviors, and so are used in more or less the same way.

[13] (Skyrms, 2000)

[14] (Matheson, 2015, p. 7)

[15] (Feldman & Conee, 1985)

1.4.3 Epistemic Peer

One of the most important concepts for the problem of peer disagreement is that of the 'epistemic peer'. Epistemic peer-hood is contrasted with epistemic superiority and inferiority. My epistemic superiors are in a better epistemic position than I am, while my inferiors occupy a worse one. Just what is meant by "better epistemic position" depends on how we measure epistemic positions, and thus how we define epistemic peers. I will refer to the relations of epistemic superiority, inferiority, and peer-hood as ones of 'relative epistemic position'. It is worth noting again that in the real world we don't only disagree with our epistemic peers. I will use to term 'interlocutor' to refer to someone who could occupy any of the relative epistemic positions. That is, an interlocutor could be my epistemic peer, superior, or inferior. Epistemic superiors and inferiors will be important actors in later chapters, but for now I will put them aside to better focus on disagreements among epistemic peers – the object of the problem of peer disagreement.

So, just what is it to be an epistemic peer, as opposed to an epistemic inferior or superior? I discuss two definitions from the literature that approach the concept from rather different angles.

Thomas Kelly tells us that:

"... two individuals are epistemic peers with respect to some question if and only if they satisfy the following two conditions:

(i) they are equals with respect to their familiarity with the evidence and arguments which bear on that question, and

(ii) they are equals with respect to general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias" [16]

Adam Elga offers his own definition:

[16] (Kelly, 2005, p. 10)

"... you count your friend as an epistemic peer with respect to an about-to-be-judged claim if and only if you think that, conditional the two of you disagreeing about the claim, the two of you are equally likely to be mistaken."[17]

Elga's definition, I think, captures both a broader and more *useful* set of possible peer combinations. If it is the case that to be "equals with respect to general epistemic virtues" peers are required to have broadly similar sets of virtues then Kelly's definition only allows for a small group of people to be considered epistemic peers, depending on what we take epistemic virtues to be. The epistemic virtues Kelly is talking about are borrowed from Gutting and include "intelligence, perspicacity, honesty, thoroughness, and other relevant virtues". [18] Interestingly, there are many ways in which people might rely on profoundly different epistemic virtues and still be epistemic peers in Elga's sense. For example, Hob and Fae might rely on different cognitive mechanisms to produce their beliefs and yet be equally likely to be right. Hob could rely on a subtle subconscious intuition where Fae is more deliberately "thoughtful". Fae might in fact be more "intelligent" than Hob, but her biases or arrogance reduce her overall odds of getting things right and make her as likely to be right as Hob is.

Kelly's first criterion also makes epistemic peer-hood a very rare phenomenon. Exactly where we draw the lines for what can and cannot count as evidence for *P* is unclear. I think that this is a big part of why it is so hard to resolve disagreements. In philosophy, for example, it seems common that when we disagree about some argument *A* we find that this is because we disagree about some other related argument *B*, which in turn is explained by another disagreement about another argument *C*. Suddenly the considerations bearing on *A* include those bearing on *C*, and so the body of evidence that supposedly (indirectly) justifies your belief about *A* grows in a fractal pattern that is (quite probably) impossible to predict by considering *A* alone. Even if we ignore this problem, could we say that Hob and Fae satisfy Kelly's first criterion on the global warming question? There are so many texts that could reasonably be counted as evidence on this question that it is very unlikely that they, or anyone else, will have encountered only the same ones. Yet, despite all these differences in

[17] (Elga, 2007, p. 485)

[18] (Gutting, 1982, p. 83)

character and evidence, by Hob's own assessment he and Fae are just as likely to be mistaken about the question of the origin of global warming. Kelly would be right to say that Hob and Fae are not peers insofar as they possess different virtues and evidence, but if they are (or think that they are) equally likely to be mistaken about *P* then I think they are peers by the metric that in a *practical* sense matters more. Why, after all, do we care about access to evidence and epistemic virtues if not because those things are indicators of how likely someone is to make correct assessments and thus have correct beliefs?

Regardless of how rare the PPD is, how we ought to respond to it also informs how we ought to respond to disagreements with epistemic superiors and inferiors (as I will discuss in sections 2.5.5 and 3.5.4). If a peer disagreeing with me has implications for my belief, then so does a disagreeing superior (and probably an inferior too). Inversely, if I think that an epistemic superior disagreeing with me is a reason to temper my belief then it seems I would need a good reason not to think a peer disagreeing with me should also affect my belief. It is easier to make sense of the other epistemic relationships (superiority and inferiority) when those relationships are determined by a single variable – as is the case with Elga's definition. The following definitions, found in Frances, illustrate this point:[19]

"If you think she is more likely … [to be correct] …, then you think she is your epistemic superior on that question; if you think she is less likely, then you think she is your epistemic inferior on that question; if you think she is about equally likely, then you think she is your epistemic peer on that question."

With Kelly's definition it is relatively easy to discount peer-hood, but it is not obvious whether that leaves my non-peer in a better or worse epistemic position than myself. If Fae is less familiar with the evidence but possesses more epistemic virtues than Hob should we say she is Hob's superior or inferior? I think that to answer this question all we need to ask is if, all things considered, we take Fae to be more or less likely than Hob to be right. In some cases, it may not be much (or any) easier to make assessments of someone's epistemic position with Elga's definition but at least once that assessment has been made it is obvious what everyone's relative standing is. There is a sense that any version of the EWV that

[19] (Frances, 2014, p. 45)

doesn't also tell us how to respond to disagreements with epistemic superiors and inferiors is incomplete and ignores the area where it could be most useful in the real world.

Given these considerations, while Kelly and Elga's definitions of epistemic peer-hood open doors to different but equally valid discussions, for the purposes of this thesis when discussing epistemic peer-hood I will be using Elga's definition.

1.5 Conclusions

With these definitions in place exactly what my statement of the problem of peer disagreement is asking should be a bit clearer. Here is the problem again:

A possesses some belief *P* based on evidence *E*. How should P be affected when *A* learns that *B*, who *A* considers an epistemic peer, disagrees about *P*?

Hob's belief is his credence of, say 0.8, that [climate change is caused by humans]. This attitude is based on the evidence he collected during his research and over the course of his life. Hob considers Fae to be his epistemic peer because Hob thinks they are equally likely to have the right credence about whether [climate change is caused by humans]. Hob and Fae disagree because they don't share the same credence that [climate change is caused by humans] (Hob's credence is 0.8, Fae's is 0.1).

So, what is Hob to do? I am looking for a practical solution to this problem. That means a solution that is designed both to address the problem as it presents itself in the real world and to be used by normal people.

Now that the problem of peer disagreement and its various components have been clearly stated I will move on to discussing how the Equal Weight View is normally taken to resolve this problem, and where it goes wrong.

***

Chapter 2. The Equal Weight View

2.1 Introduction

Philosophers have attempted to solve the problem of peer disagreement (PPD) with a diverse set of solutions,[20] but my focus is on the Equal Weight View (EWV). Because I am interested in developing a practical application of EWV, rather than arguing about whether or not it is ultimately correct, I won't be doing much to justify the position or to defend it from hostile arguments. I will assume that the underlying principles of the EWV are correct and that we should be using those principles to respond to the PPD. As I don't intend to offer a defense of the EWV, I will also avoid explaining and reviewing the host of solutions competing with it. One solution, the Steadfast View (SV), originally defended by Kelly,[21] is an exception because I will argue that something like it has a surprisingly important part to play in the Idealized Equal Weight View (IEWV) in Chapter 3. As Frances and Matheson point out, the SV is the view that most dramatically opposes the EWV.[22]

When discussing what a practical solution to the PPD would look like in section 1.3, I discussed how a solution could run into internal and external practical problems. In this chapter I will discuss both types of problem, but my focus will be on the external problems with the EWV. That is, the ways in which even if it were applied perfectly, the standard EWV would fail to achieve its goals (and violate its own principles) when faced with real-world disagreements.

Thus, in section 2.2 of this chapter I briefly outline how the EWV and SV respond to the problem of peer disagreement. In section 2.3 I discuss the underlying three principles of the EWV and explain how they are normally applied to cases of disagreement. In section 2.4 I make a distinction between psychological and non-psychological evidence. This distinction is useful for explaining where the EWV and SV diverge and will be used extensively in Chapter 3 where I develop the IEWV. In section 2.5 I compare the EWV and SV by discussing scenarios from the literature that intuitively favor both views. In section 2.6 I discuss a set of

[20] Beyond the EWV, see Thomas Kelly's initial Steadfast View (2005), his later Total Evidence View (2010), and Jenifer Lackey's Justificationist View (2010).

[21] (Kelly, 2005)

[22] (Frances & Matheson, 2019)

problems that would arise during real-world use of the standard EWV. These are the *external* practical problems with the EWV. Again – addressing these problems is the primary goal of the IEWV.

2.2 Outlining Solutions to the Problem

Before moving forward, I want to quickly outline the two solutions to the PPD that I will discuss in this thesis, the EWV and the SV. While I will not argue that either of these positions is the correct solution to the PPD, I do want to consider what is attractive about them when they are put into practice. This is so that in Chapter 3, where I develop the Idealized Equal Weight View, it is clear what that application needs to do well and where it might be better than standard applications of the EWV. So, starting with the EWV, let's quickly consider how these two theories propose we solve the PPD.

2.2.1 The Equal Weight View

The EWV holds that disagreeing with an epistemic peer gives me a very good reason to revise my belief. This revision is achieved by conciliating my belief with my peers. Assuming, as I have, that we can use credences to meaningfully describe beliefs, I will conciliate my belief with my peers by, what David Christensen calls, "splitting the difference" between our beliefs.[23] I am using credences to talk about beliefs, so I need to split the difference between our credences that *P.* If my credence that *P* is 0.3 and my peer's is 0.5, then the EWV will recommend I adopt a credence of 0.4.

2.2.2 The Steadfast View

Unlike the EWV, the SV holds that the PPD is not at all threatening to the rationality of our beliefs, so we should remain 'steadfast' in our beliefs, making no attempt to revise them merely on the grounds that we disagree with our peers.[24] What is rational for me to believe is a function of how *I* interpret the evidence, not of how my peers have interpreted it.

[23] (Christensen, 2007, p. 203)

[24] (Kelly, 2005)

2.3 The Principles of the Equal Weight View

The EWV argues for a conciliatory approach to peer disagreement based on the following three foundational principles (adapted from Frances and Matheson):

"Defeat: Learning that a peer disagrees with you about *P* gives you a reason to believe you are mistaken about P.

Equal Weight: The reason to think you are mistaken about *P* coming from your peer's opinion about *P* is just as strong as the reason to think you are correct about *P* coming from your opinion about *P*.

Independence: Reasons to discount your peer's opinion about *P* must be independent of the disagreement itself." [25]

For now, I will work through Hob's disagreement with Fae to demonstrate these principles in practice in the standard application of the EWV.

Hob considers Fae to be his epistemic peer about *P* ([global warming is caused by humans]). Because of the principle of defeat their disagreement about *P* gives Hob a reason to think that he might be mistaken about *P*. The principle of independence means that Hob can't use his disagreement with Fae about *P* as evidence that Fae is not his epistemic peer regarding *P*. So, Hob, having thought of Fae as his peer going into the disagreement, has no new reason to disregard Fae as his peer. The principle of equal weight means that the weight of Hob's own assessment and the assessment of his epistemic peer are doxastically equal. Thus, Hob cannot reasonably favor his credence that *P* over Fae's. He must assign equal weight to the two credences.

Hob should become less confident in his belief that global warming is caused by humans when he learns that his epistemic peer, Fae, disagrees with him. Because Hob and Fae both think they are equally likely to be right, they ought to afford each other's beliefs equal confidence. So, if Hob thinks *P* is true and Fae thinks *P* is false, they should both be agnostic

[25] (Frances & Matheson, 2019)

about *P*. Equal weight will of course be given to more granular attitudes about the truth of *P*. If Hob has a credence of 0.8 in *P* and Fae has a credence of 0.1, equally weighting their views would produce a credence of 0.45 that *P*. Thus, Hob should go from being fairly sure that *P* is true to being almost agnostic about *P*, leaning ever so slightly towards thinking that it is false.

2.4 Psychological Evidence

Much of the debate surrounding the PPD comes down to the role of evidence. A distinction is drawn between two types of evidence. Understanding this distinction is crucial for understanding how the EWV should be put into practice. I will use the distinction of psychological evidence (PE) and non-psychological evidence (NE).[26] Gardiner states that non-psychological evidence "bears directly on whether [*P*] without concerning other people's judgments".[27] Psychological evidence, he says, concerns "what other people believe regarding whether *P*".[28] PE *only* includes evidence derived from what *other people think*.

Part of Hob's NE for his belief that climate change is caused by humans is his understanding of how greenhouses gases affect climate. My NE for my belief that it is a sunny day outside is the clear sky I see through my window, the warm sensation that I feel as sunlight hits me, and my reasoning that these phenomena are indicative of a sunny day. NE is what we are usually talking about when we discuss evidence (e.g., "bloodstains are *evidence* of a murder"). For Hob, the fact that he disagrees with Fae, his epistemic peer, is PE that either his assessment of the NE was wrong or that his NE was itself not particularly good and thus led to a false belief. That he has read that scientists interpret *their* NE as indicating that global warming is caused by humans is also PE for Hob. That you and I agree that it is sunny outside is, for me, PE that I interpreted my NE (my perceptions of a clear sky, warm sensations, and corresponding reasoning) correctly. When we rely on the testimony of others

[26] In the literature it is more common to find a distinction drawn between first-order and higher-order evidence (e.g. in Kelly (2010)). The distinction similar, with PE playing the role of higher-order evidence and NE that of first-order evidence. However, the first/higher-order distinction does not specific the source of the evidence. I favor the PE/NE distinction here because specifying those sources it makes it easier to see exactly where the standard EWV makes some of its mistakes.

[27] (Gardiner, 2014, p. 11)

[28] *Ibid*.

we are relying on PE. Often relying on testimony, especially that of experts, means trusting PE when little or no NE pertaining to the question at hand is available to us (or where we lack the necessary skills to interpret that evidence).

The distinction between PE and NE is useful for drawing a clear line between the Equal Weight View and Steadfast View, as well as understanding why both views can be attractive.

In the EWV PE is very important and can end up overruling NE. That is to say, *A*'s belief that *P* (based on NE), is subject to change when PE about *A*'s belief is introduced. For Hob this means lowering his confidence in his belief *P* (that [global warming is caused by humans]) by assigning equal weight to his and Fae's positions when he learns that she disagrees with him about *P*.

In the SV, NE is the sole determiner of the rationality of a belief. What would be PE is, for reasons outside of the scope of this thesis, taken not to bear on the rationality of one's beliefs. Thus, according to the SV, whether it is rational for Hob to believe *P* with the conviction that he does is determined purely by Hob's own analysis of his non-psychological evidence. That someone Hob takes to be his epistemic peer disagrees with him about *P* does not mean that he should reduce his original credence in *P*.

Without getting distracted by arguments about which position is correct, I think that it is fairly intuitive that PE should in some way factor into what we ought to believe. If for no other reason, it seems to me that what other people think about *P* could be useful information if I don't completely ignore it. Even if I weren't to give as much weight to PE as the EWV recommends, it *seems* like it would be wasteful and unreasonable to put it aside entirely, as the SV holds we should.

2.5 Comparing the EWV and SV

While I have assumed that the EWV's sensitivity to PE is correct, I do think there are areas where the SV appears to be more attractive. I want to highlight some cases from the literature that seem to favor each position now. My intention in this chapter is not to justify the EWV, but to highlight some areas where might fall short. The SV appears to lead to better outcomes in some of these cases, but (assuming the principles of the EWV) it comes with its own

problems. In Chapter 3, where I develop the IEWV I argue that we can have the 'best of both worlds' while also avoiding their major shortcomings.

2.5.1 Radical Progress

There is something about the epistemic path that the SV demands that I find, in the right light, intuitively admirable. A willingness to uncompromisingly stand by a belief regardless of how convinced epistemic peers (and indeed superiors) might be that said belief is wrong has been necessary for great historical advancements in so many domains of human inquiry. We see this clearly in examples of scientific breakthroughs like heliocentrism or evolution, or in moral and political advancements like the anti-slavery or global enfranchisement movements. I think it is likely that many of the radical ideas that have improved our understanding of the world would not have been possible without their originator's steadfast commitment to their own understanding of the NE despite disagreement with their epistemic peers and superiors. The SV is appealing precisely because it mandates the kind of intellectual individualism that makes these advancements possible.[29] Because we rely only on our reading of the NE, we are each ultimately responsible for the content of our own minds. This creates the space needed for big changes in how we think about the world. I will examine why the standard EWV would be much worse at enabling epistemic progress in section 2.6.2 where I discuss the election magnetism problem.

2.5.2 The Restaurant Case

There are cases where the EWV's sensitivity to PE seems to fare better than the SV's unyielding commitment to NE. Christensen's Restaurant Case is commonly deployed as a thought experiment that suggests something like the EWV must be correct.[30] The problem unfolds as follows:

A group of friends are ready to pay the bill after a meal at a restaurant. The friends have agreed to split the bill evenly and to add a 20% tip. They are handed the bill and after taking a minute to do the mental math one friend, whom I'll call *A*, determines that they each owe $43, while another (*B*) concludes that the amount should be $45. In this case *A*'s NE is her

[29] (Weisberg & Muldoon, 2009, p. 244)

[30] (Christensen, 2007, p. 17)

understanding of mathematics and the processes she deployed to arrive at $43. Her PE is that her peer, *B*, disagrees with her.

Assuming that this is a case of peer disagreement (the two friends take each other to be equally reliable bill calculators based on their "track records"[31]) it seems fairly intuitive that they ought to follow the EWV and assign equal credences to both positions ('n=43 'and 'n=45'),[32] at least initially and until they have carefully checked their solutions. *A* recognizes that while it was a relatively simple procedure, the process she used to solve this math problem is vulnerable to small mistakes which are not immediately detectable. This is just a basic background acknowledgment of her fallibility as a thinker – one that we ought to share. Surely, without a reason to think that it was her peer rather than herself who made the mistake, it would be reckless for *A* to believe in these circumstances that she is right and that her epistemic peer *B* is wrong. Yet it seems that the SV maintains that the PE here is toothless and would demand that *A* ignore it. In this case the steadfast independence that SV calls for looks more like hubris than integrity.

2.5.3 The Extreme Restaurant Case

But there are also scenarios where sticking to your guns as the SV prescribes aligns better with our intuitions than the conciliatory approach of the EWV. Christensen's *Extreme* Restaurant Case[33] is one. The scenario plays out in the same way as the case above, but *B* thinks that the amount each of them owes is $450, a value higher than the total bill for the table (*A* still calculates the bill to be $43). This thought experiment is starting to stray into the realm of untethered philosophical idealization, but we do, however rarely, deal with people who have made mistakes as dire as the one *B* made here. Here it seems that *A* would be much better off being steadfast and ignoring her peer's obviously confused position, peer or not. There is contention in the literature about whether this case shows that the principle of independence is false because it appears *A* is forced to use *B*'s opinion about the disagreed

[31] *Ibid*.

[32] Note that this means thinking that [n=43] is as likely to be true as [n=45], not melting the beliefs together to arrive at [n=44].

[33] (Christensen, 2007, p. 17)

upon belief to say that *B* is not her peer.[34] The EWV might be amended with a clause that allows it to handle situations where a mistake is obvious and undeniable - though such a fix seems ad-hoc to me. I think that the EWV actually handles this case just fine. The EWV only holds that the PE coming from the fact of peer disagreement is a reason for *A* to shift her credence that *P* towards her peers'. By the principle of Independence, *A* can't use *B*'s position in the disagreement as a reason to discount his opinion, but evidence external to the disagreement can be used to do just that.[35] *A's* belief that a part cannot be larger than the whole *it is a part of* is such external evidence. So, the scenario plays out like this: *A* and *B* disagree which serves as PE that *A* might be wrong about *P*. *A* ought to follow the EWV and conciliate their positions. However, if *A* thinks at all about *B*'s position, she will see that it violates a basic principle she is very confident in. That is, *A* has (external) NE that *B* must be wrong, and thus that she should not consider *B* her peer on this question (because *B must* be wrong), and thus that the PE does not suggest that they are equally likely to be wrong. I think that this is representative of the process that many people would follow in this case. When I learn that my peer thinks that the amount owed is $450, I do have, for a moment, a reason to think that I could be wrong. It is just that I will almost immediately realize that they must have made a mistake (perhaps before I can even perform the operation of tempering my credence). If for some reason I fail to realize that *B*'s position must be wrong then I should follow the EWV as normal. That is, contrary to the SV, that my peer says I owe $450 *does* give me a reason to doubt my belief, but that reason is fragile and can be quickly defeated by other evidence. So, we don't need to subscribe to the SV to avoid being swayed by obviously wrong positions, even if they belong to our epistemic peers.

2.5.4 Bootstrapping

The process of using external evidence to reassess someone's peer-hood outlined above should not be confused with what Elga calls "bootstrapping".[36] Elga argues that an important point in favor of the EWV is that it does not lead to "bootstrapping". The bootstrapping problem occurs when one wishes to lower the weight they should assign to their peer's position to less than the weight they assign to their own. By definitions being used in this

[34] According to (Frances & Matheson, 2019), see (Lackey, 2010).

[35] (Frances & Matheson, 2019)

[36] (Feldman, 2007)

thesis we know that in the PPD we must take our peers to be as likely as we are to be right about *P*. Thus, I will need a reason for assigning less than equal weight to my peer. Elga holds that having such a reason would require that somewhere along the line I acquire new evidence that shows my peer is less likely to be right about *P*. If there is no such evidence, then I am using the outcome of my method of thinking (an attitude to *P* is in conflict with my peers') to bolster my confidence in said method (whatever processes I used to determine if *P*). This, Elga says, is akin to using the products of my visual system to verify that my visual system is working.[37] That I can't do so is an important part of Feldman's Dean in the Quad thought experiment.[38]

In this scenario, two people, who I will again call *A* and *B,* are staring out of a window into a quad. Upon speaking they realize that *A* thinks she sees the dean standing in the quad while *B* thinks the quad is empty. They take themselves to be peers regarding their faculties of vision and think that this is a genuine case of disagreement (it is clear that neither is joking). Here the extent of the 'mistake' is a fair bit greater than in the Restaurant Cases; it looks like someone is hallucinating or suffering from some other serious perceptual malfunction. But there appears to be no reason for either party to assume that it was their peer rather than themselves who made that mistake. Unlike in the Extreme Restaurant Case where *A* has external evidence that *B* is wrong, *A* has no reason to think that *B*'s faculties of vision are malfunctioning rather than her own. Here the PE does forcefully suggest that either *A* or *B* is mistaken, and there is no other evidence discrediting that PE. For A to follow the SV here and favor her own belief would make her guilty of bootstrapping.

2.5.5 Epistemic Superiors and Inferiors

When comparing the EWV and SV it is also worth looking at how they handle cases adjacent to the PPD. Often enough our disagreements are not only with our epistemic peers. We are also surrounded by epistemic superiors and inferiors who are certain to disagree with us on many issues. As the EWV calls for an equal weighting of the views of epistemic peers it is quite sensible that it should also end up calling for an *unequal* weighting of the views of superiors and inferiors, where we grant less epistemic weight to our inferiors and more to our

[37] (Elga, 2007, p. 488)

[38] (Feldman, 2007)

superiors. Just how slanted this weighting should be depends on just how big the epistemic gap between the parties is. The more someone is your superior the more weight their positions ought to be given. This all seems quite intuitive to me. Generally, we want to be able to trust what people we take to be greater experts on matters than ourselves are saying while devaluing the opinions of the uninformed. Thus, compared to how we weight our own beliefs, we assign greater epistemic weight to the beliefs of superiors and less epistemic weight to the beliefs of inferiors. Does it not then follow that we should assign equal weight to our epistemic peers? I take it that the EWV handling - or being contiguous with solutions that handle - other forms of disagreement in sensible ways is a point in favor of the solution.

What does the SV have to say about disagreement with epistemic superiors? Either it maintains the primacy of NE or concedes that PE has some role to play. If it maintains that NE is the only evidence that matters, then the fact that *A* disagrees with *B* on *P* where *A* thinks *B* is *more likely than her* to be right about *P* is no reason for *A* to doubt her position on *P*. If *A*'s goal is simply to put herself in the position with the best chances of being right about *P* it seems obvious that the views of someone she considers *more likely than herself* to be right about *P* should have some bearing on her belief. Alternatively, if the SV says that in cases of disagreement with epistemic superiors we ought to be sensitive to the PE stemming from our disagreement then we need a reason for why we have nothing to worry about when our disagreement is with people who are *merely* epistemic peers rather than superiors. Again, if there is utility to be found in being sensitive to the views of epistemic superiors it seems that there must also be some utility in being sensitive to the views of epistemic peers. We would need a good reason for approaching our disagreements with epistemic superiors and peers with radically different strategies.

2.6 Problems with Actual Applications

Usually, when the PPD is discussed it is in idealized thought experiments like the Dean in the Quad or Restaurant Case. In these thought experiments we are dealing in cases of peer disagreement that involve only two peers. In reality, we are often aware of many epistemic peers (if not also superiors and inferiors) who disagree with us. If we attempt to use the EWV to make sense of disagreements that involve more than two people, as we would in real-world applications, a few new quirks of the theory become apparent. Some of these quirks lead to

serious problems where the beliefs produced by applying the EWV don't appear to represent its principles. These problems will be addressed by the IEWV in Chapter 3.

2.6.1 The Skepticism Problem

A feature of the EWV is that the beliefs it produces are the product of "splitting the difference" with our epistemic peers.[39] Often enough this will mean adopting a position that is more agnostic towards the belief in question. If you believe *P* and I believe ~*P* the middle ground is agnosticism about *P*. Thus follows Feldman's worry that this kind of conciliation leads to skepticism.[40]

If we are measuring degrees of belief with credences, then a credence of 0.5 in *P* is agnosticism about *P*. It would be rare to arrive at a credence of exactly 0.5. As Carey and Matheson point out, if enough beliefs are counted the EWV will rarely result in a tie.[41] However, I think it is fair to say that there is a range of credences close to 0.5 that are so weakly inclined towards believing or disbelieving *P* that they are *practically* indistinguishable from agnosticism.

For most people there are some cases where moving towards agnosticism will violate intuition or otherwise be viewed as a most unpalatable option. There are undoubtedly people we ought to consider our epistemic peers on the other side of just about every issue we are likely to care about. This is true whether the debate is political, religious, or philosophical in nature (indeed, this is especially the case in philosophy).

So, the EWV will often lead us to agnosticism or something close to it about many of our beliefs. This will not mean adopting general skepticism, but for many beliefs the result will be difficult to distinguish from skepticism. For many people this is a most unwelcome result.

On the other hand, most people seem to find themselves swung one way or another by the NE on these questions. Given that NE analysis, for whatever reason, usually results in a more

[39] (Christensen, 2007, p. 203)

[40] (Feldman, 2007)

[41] (Carey & Matheson, 2013, p. 17)

assured attitude towards *P*, the SV is unlikely to lead to pervasive agnosticism about our beliefs.

When dealing with granular degrees of belief it is more likely that we will end up slightly favoring one position over another, rather than arriving at a perfectly agnostic credence of 0.5. Still, a credence of 0.4 or 0.6 is a relatively weak level of conviction, and we should expect that, in some cases where we really care about our beliefs, the EWV will take our beliefs straight into this range of murky uncertainty. A slight inclination towards one side might be enough to decide how to act in a pinch, but it is unlikely to satisfy people made uncomfortable by uncertainty. I don't think this is something that any application of the EWV making use of credences can avoid.

2.6.2 The Election Magnetism Problem

Just how far the EWV drives us towards or away from agnosticism about a specific question depends on what our epistemic peers think about that question, and thus what our PE looks like. When we encounter and disagree with many peers we must process the views of many peers. Carey and Matheson describe the act of applying the EWV to this group of peers as looking like an "epistemic election", the result of which we should adopt as our own belief.[42] This is a particularly useful way to conceptualize the EWV. When I apply the EWV by assigning equal weight to all of my peers, relatively little weight is left for me to allocate to my own position. While my 'vote' counts, it is just a drop in the bucket. Thus, the belief-forming process the EWV suggests is hardly different from using an election to determine what we ought to believe when enough people are involved. Now, elections are very useful tools, but if a population relies solely on an election to determine what they should believe they will run into a serious problem.

Imagine a political election where everyone decided who to vote for based on who *won* the last election. In such a scenario the election would stop being a useful tool. It would get stuck in an input-output loop where we vote for party *X* because they win elections, and because we vote for them, party *X* keeps winning elections, and we continue to vote for them because

[42] (Carey & Matheson, 2013, p. 10)

they win elections. I will call this 'the election magnetism problem', because we ourselves drawn closer and closer to the output as if it were attracting us like magnets.

Here we see one way that the EWV could remove the link between what a population believes and them actually *thinking* about NE. The analogous epistemic case is one where after applying the EWV everyone in a population converges on a credence of 0.4 that *P*. At this point their beliefs are untethered from their own analysis of the NE - they all hold a credence that 0.4 after using the EWV, and because they all hold a credence of 0.4, successive applications of the EWV will continue to tell them to hold a credence of 0.4. This kind of epistemic stability and consensus is just the result of the PE-focused system coming to a natural equilibrium. However, this equilibrium also means that there is no room for any of the revolutionary epistemic advancements I discussed in section 2.5.1. Never mind radical change, if everyone just adopts the election results as their beliefs it becomes hard to see how any of our views could change over time. If my primary belief generation tool is just to aggregate PE, I don't need to spend time and energy analyzing NE on my own. If I do decide to analyze the NE and disagree with the aggregate, I will find that my new beliefs cannot gain traction with my peers, because they are simply reproducing the election results, and those results are unlikely to depend on my input. If everyone follows the EWV the doxastic magnetism of the consensus is too strong to be overcome by a good idea. Indeed, as the would-be maverick theorist re-applies the EWV to resolve her newfound disagreements, she will quickly find that her own beliefs drift back towards, and are replaced by, the shared consensus. Thus, one problem that the IEWV needs a way to overcome is that of the epistemic stagnation which results from the wholesale adoption of the epistemic election results.

2.6.3 The Sequence Problem

Because the EWV is usually taken as a solution to disagreement between two peers the math involved in applying the view is relatively simple: we add up the two credences for *P* then divide by 2 to arrive at the conciliatory level of credence the EWV recommends. Let's say that *A*'s original credence that *P* was 0.4, and after disagreeing with *B* (who has a credence of 0.8 that *P*) the EWV recommends adopting a credence of 0.6 ($\frac{0.4+0.8}{2}$). *A*, following the EWV, adopts a credence of 0.6 in *P*. So far there aren't any problems; that is fairly simple math.

However, as Gardiner points out, problems start to arise when we introduce another peer later.[43] *After* applying the EWV to her disagreement with *B*, *A* meets another peer, *C,* who also disagrees about *P*. Again, applying the EWV, *A* adds her *new* credence of 0.6 and *C*'s credence of 0.2 to arrive at 0.4 ($\frac{0.6+0.2}{2}$). Once again *A* only needs to do some simple math. However, something is amiss. If *A* had encountered *C* before *B* her final credence that *P* would be 0.55 rather than 0.4 (0.3 after disagreeing with *C* ($\frac{0.4+0.2}{2}$), and 0.55 after disagreeing with *B* ($\frac{0.3+0.8}{2}$). This means that the EWV appears to produce different results depending on the *sequence* in which we encounter peers, and thus PE. In this case, the difference in credence is not substantial (0.15), but it is easy to see that with a long enough series of peers (or with enough variance in those peers' beliefs[44]) a different sequence of encounters could result in radically different positions being recommended by the EWV. This can't be right; surely the order that *A* encounters peers in ought to be irrelevant to what she believes about *P?* Indeed, Gardiner argues that this would violate the commutativity of evidence principle, which states that:

"To the extent that what it is reasonable for one to believe depends on one's total evidence, historical facts about the order in which that evidence is acquired make no difference to what it is reasonable for one to believe."[45]

I will refer to this issue as the sequence problem.

2.6.4 The Stubborn Peer Problem

Not all peers that *A* disagrees with are also going to be proponents of the EWV. What happens if *B* is an adherent of the SV, or is otherwise unwilling to conciliate their belief? The

[43] (Gardiner, 2014, p. 5)

[44] Note that the order we encounter peers matters for the effective weighting of peers who *agree* with us (i.e. share our credence that *P*). Say I have five peers. I and four of these peers have a credence of 0.3, while the fifth has a credence of 0.9. If I meet the four peers who agree with me before the fifth, it will be as if I hadn't met them at all. But if I meet the fifth peer first the other four will pull me back towards 0.3. Clearly, however we solve this problem we will need to find a way to count peers who agree with us.

[45] (Gabbay et al., 2011, p. 451)

following problem is discussed in detail by Brian Weatherson.[46] When *A* first encounters *B* and disagrees with him she will apply the EWV. Once again, we will say that *A* and *B*'s credences are 0.4 and 0.8 respectively. So, after the initial disagreement *A* comes to have a credence of 0.6 ($\frac{0.4+0.8}{2}$). But *B* remains steadfast, so his credence remains at 0.8. Some time later, *A* encounters *B* again. So, *A* once again finds herself in disagreement with *B*. Because *A* is an adherent of the EWV she will respond to the disagreement by applying the EWV. After doing so she will arrive at a credence of 0.7 ($\frac{0.6+0.8}{2}$), and *B* will again remain steadfast in his credence of 0.8. This process will repeat itself, and *A* will find herself shifting asymptotically towards *B*'s position. Again, it looks like something has gone wrong. How epistemically stubborn *A*'s peers are should be irrelevant to what she ought to believe about *P*, but on this reading of the EWV it looks like it is a factor that can override all other considerations.

Even in the best cases of EWV adoption, it is very unlikely that a whole population could be convinced to be proponents of the EWV. The population will include people who subscribe to the SV or something like it, or any one of the other solutions to the PPD. Assuming that most of the population follows the EWV we will see that this minority of SV adherents is likely to have an outsized influence on the beliefs the EWV group converges around. Because we are dealing with a whole population of EWV adherents there are many more vectors through which the stubborn peer can influence the EWV process. Each EWV adherent that the stubborn peer meets must factor in the stubborn peer's belief in their assessment. Their position thus shifts slightly towards the stubborn peer's, and when they present their beliefs to other EWV adherents they will present this slightly shifted belief. If the stubborn peer can influence enough of the population's EWV adherents at once he can permanently shift the consensus towards his own position (due to the result magnetism problem outlined above). The stubborn peer makes no revisions on his beliefs, so he (and his compatriot SV adherents) will constantly exert this pressure to revise on the EWV population. The entire population could thus end up consolidating around beliefs that are in no way representative of an equal weighting of all their original positions. We end up not only with stagnation, but a stagnant consensus that doesn't even reflect the population's original NE assessments.

[46] (Weatherson, 2013, p. 10)

Part of the reason that the stubborn peer might gain too much influence over the epistemic election is that when everyone is using the EWV it becomes incredibly difficult to know if a peer's beliefs have been accounted for.[47] Say *A* has met and applied the EWV to disagreements with *B*, *C*, and *D*. Then she meets *E* and wants to apply the EWV to his belief too. But *E* is also an adherent of the EWV. If *E* has also met *B* and applied the EWV to their disagreement, then when *A* tries to use the EWV in her disagreement with *E* she will be counting *B* twice: once directly as part of her application, and a second time when she adds *E*'s credence (because it is in part composed of *B*'s credence). Assuming she assigned equal weight to *A*, *B*, *C*, *D*, and *E* in the application, she would unknowingly be assigning slightly more than equal weight to *B*, and slightly less to *E*. A related problem is that while *A* knows her belief is the product of her and three peers' beliefs, she does not know how many peers *E* has met. Depending on this number she could easily over or underweight his position. If *E*'s belief is the product of applying the EWV to a total of 20 peers, as opposed to *A*'s four, *E*'s position should be granted much greater weight than *A*'s. But even if *A* knows that *E*'s position is comprised of 20 peers, she might not know how good *E* is at assessing peer-hood, and so may be unsure of the quality of his PE assessment.

These problems arise because of how the peers using the EWV are generating their beliefs. The EWV is, I argue, only useful when applied to beliefs that are the product of an analysis of NE – the only kind of beliefs that the SV produces. After all, it is because someone as likely to be right about *P* as I am has *thought about P* that, by the principle of defeat, their disagreeing with me indicates that I might have made a mistake. If my peer has not *actually* thought about *P* then their position on *P* is irrelevant to me.[48] Thus, if my peers thinking about *P* (i.e., their interpretation of the NE which led to their original (unconciliated) credence in *P*) is only a small part of their current credence in *P* (which they arrived at after applying the EWV) then their current position is only threatening to mine to the (now limited) extent that it is the product of their NE.

[47] (Gardiner, 2014, p. 10)

[48] (Gardiner, 2014, p. 6)

2.7 Conclusions

In section 2.5 I discussed scenarios that highlight boons of both standard applications of the EWV and the SV. In some scenarios the EWV is an intuitively attractive solution, while in others the SV seems better. When it comes to substantial, paradigm-shifting epistemic change, the EWV appears to leave us without the epistemic independence required to have radically different and unpopular beliefs. Yet, the restaurant case and charge of bootstrapping suggest that the independent thinking of the SV can stray into something that looks a lot like dogmatism. This is especially the case when we extend the PPD to cases that also involve epistemic superiors disagreeing with us. In this thesis I have assumed that the principles of the EWV are correct. I have discussed areas where the SV seems to outshine the EWV because I think that the IEWV I develop in Chapter 3 allows us to follow those principles while also adopting some of the attractive features of the SV in these areas.

In section 2.6 I focused on scenarios where the standard application of the EWV fails to put its principles into practice, highlighting a number of external practical problems with the view. These scenarios, rather than idealized isolated cases of disagreement, are designed to be representative of the persistent and recurring disagreements we encounter in the real world. The problems are that the EWV seems to lead to too much agnosticism, produces a static epistemic landscape, recommends different beliefs depending on the order we encounter peers in, and doesn't work when applied to disagreements with SV adherents. I think that the defining feature of these problems is that they involve more than two peers whose beliefs need to be processed[49]. Most of these problems are results of the mechanisms of the EWV that do that processing. The EWV is being applied to sets of beliefs that are, in part or in full, the result of PE analysis. That is, the beliefs serving as input for an application of the EWV are themselves the product of an application of the EWV. In Chapter 3 I will argue that we should only use beliefs that are the product of NE analysis as inputs in the IEWV. To some extent I think that the steps required to rebuild the EWV to do this will even protect it from the charge of leading to too many agnostic beliefs.

[49] I will argue in section 3.2 that the stubborn peer mechanically occupies the position of multiple peers.

Now, with a firm grasp on where the EWV falls short as a solution to the PPD in the real world, I will argue for a revised application.

***

Chapter 3. An Idealized Equal Weight View

3.1 Introduction

I will now present a version of the EWV, which I call the Idealized Equal Weight View (IEWV) that overcomes the external practical problems outlined in section 2.6. At the same time, I think that the IEWV borrows from the SV in ways that allow it to keep the most attractive parts of both the SV and EWV.

Recall from section 1.3 that it is important for a practical solution to account for the limits of whoever is meant to use it. While it is idealistic in some ways, the IEWV is not idealized by way of ignoring these limits. I assume all the human limitations that might get in the way of applying the EWV are still present for anyone attempting to apply the IEWV. However, while the application never demands the impossible by asking us to overcome these hard limitations, it does assume a willingness to spend more resources to run the EWV than is reasonable for broad or persistent adoption. Instead of being impossibly difficult, the IEWV is, I think, unrealistically tedious. Thus, the idealized application is analogous to the computer-generated marathon training plan that never asks more of us than we are capable of giving but is nonetheless too demanding in other areas. If *all* we cared about was applying the EWV, the IEWV would be a good practical solution. At least, it is something we *could* do.

So, in developing the IEWV I don't intend to ignore the limitations that would apply to a real-world application of the EWV, but to see what the best case for an application within those boundaries would look like. That best case, being so resource expensive relative to our limits, will be too demanding for broad adoption, but it can serve as a schematic for what the ideal processes involved in an application of the EWV are. The IEWV is sensitive to the real-world environment it is intended to be used in, and thus overcomes the problems discussed in section 2.6. It is just less sensitive to what would make it the best solution for its real-world user.

Later, in Chapter 4, I will use the application model developed here to attempt to provide more cost-effective substitutes for those processes, hopefully arriving at a reasonably efficient and actionable application: the Practical Equal Weight View (PEWV).

In section 3.2 I explore how the order of operations in the standard EWV is the source of many of its shortcomings. We should instead adopt a parallel processing model. In section 3.3 I explore functional problems with parallel processing that will inform how the IEWV is organized. In section 3.4 I outline the primary differences between the IEWV and the standard approach. In section 3.5 I describe the contents of the ledger, a key feature of the IEWV. Finally, in section 3.6 I describe just how the IEWV is used.

3.2 Making the EWV work in actual applications

Before I can explain the complete mechanics of the IEWV, I need to provide solutions to the practical application problems discussed in section 2.5. How we solve these problems will inform the design of the IEWV.

The sequence and stubborn peer problem share a solution because they are the result of the same mistake. In both cases the EWV is being applied in a flawed way; by a method that does not actually reflect one of the underlying principles of the EWV: equal weight.

Say *A*, *B*, and *C* are all epistemic peers,[50] so there is no reason to think that any one of them is more likely than the others to be right about *P*. Counting her own and her peers' positions, *A* has collected three credences belonging to three peers. They are peers, so she should assign equal weight to these credences. It looks like each of their credences should account for 33.3% of the credence the EWV recommends.

However, as is noted by Gardiner,[51] a problem occurs because *A* has applied the EWV in series; first to her disagreement with *B*, and then to her disagreement with *C*. Thus, at the end of the sequence *A* and *B*'s credences have effectively received half of the weight of *C*'s. As they constitute the final result, *C* is weighted 50%, but *A* and *B*'s original credences are only being weighted at 25% each. I will call this method of weighting views 'series processing'.

[50] Or are considered to be epistemic peers by *A*.

[51] (Gardiner, 2014, p. 10)

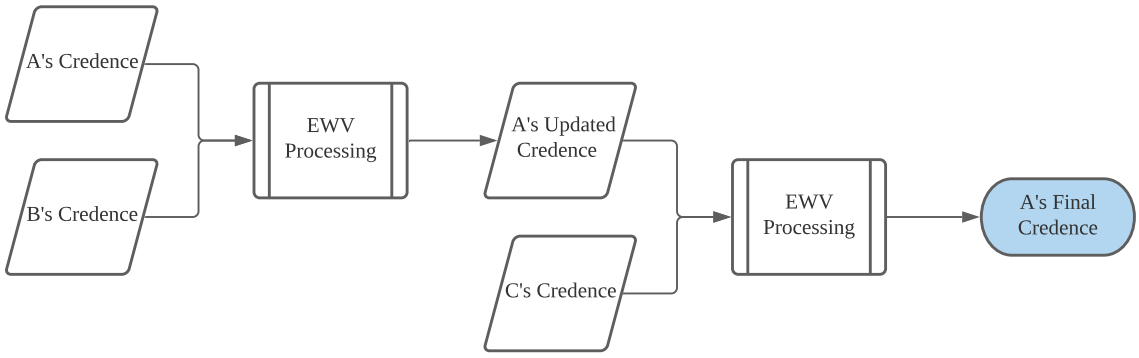

Fig. 1) Series Processing

This unbalanced weighting occurs because *A* and *B*'s credences each account for half of *A*'s new credence after she encounters *B*, and thus only half of the credence that *A* brings to her disagreement with *C*. That is, after disagreeing with *C, A*'s credence is $\frac{\left(\frac{A+B}{2}\right)+C}{2}$. If the three credences were actually assigned equal weight, A's credence would instead be $\frac{A+B+C}{3}$.

I think that the same serialized processing of credences is what allows *A* to be swayed towards *B*'s position in the stubborn peer problem (where *B* is the stubborn peer). Each and every time *A* applies the EWV to her disagreement with *B, B* is effectively playing the part of another peer in the series, and so has his credence systematically assigned a greater total weight than *A*'s. Upon each recursive application of the EWV with *B* this disparity grows. At the second instance of disagreement *B*'s views are weighted as if he were *B and C* in the sequence problem, in the third as if he were *B, C, and D*.

So, in a series processing application, every subsequent peer has their belief about *P* weighted slightly higher in proportion to the peers encountered before them. Let's say *A* always applies the EWV when she meets a peer with a belief about *P*. After meeting only nine peers the credence of the first peer *A* met in the sequence has an effective weight of less than 0.2% in the calculation of the recommended credence in *P*, while the ninth peer's credence (*J*), being the one most recently encountered, is weighted at 50%. If these epistemic peers really had their beliefs assigned equal weight *each* of the 10 peers' credence would account for 10% of the recommended credence – as would be the case if *A* met and disagreed

with all of these peers at once (e.g., at a seminar instead of in series over the course of her career).

Fig. 2) Equal Weighting, 10 Peers

Fig. 3) Series Processing Weighting, 10 Peers

So, the problem is that the series process, where I conciliate two beliefs at a time and repeat the process, can only afford equal weight to the two credences in its most recent application, rather than to all the credences that are constitutive of the final recommendation if more than two people have disagreed. Gardiner discusses a number of solutions to this problem, but the most promising is simply to apply the EWV, so to speak, in parallel.[52] That is, to process all the input credences at once using the $\frac{A+B+C}{3}$ model. The parallel processing model will always assign equal weight to the credences I have collected because all of those credences are direct inputs every time the process is used.

[52] (Gardiner, 2014, p. 11)

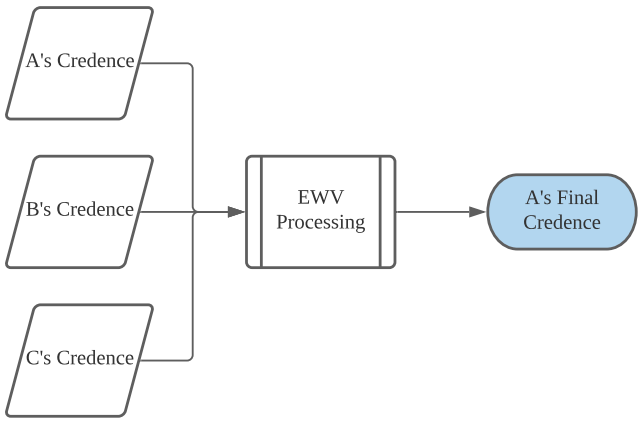

Fig. 4) Parallel Processing

3.3 Difficulties with the Parallel Process

Gardiner points out that using a parallel processing application of the EWV would require me to remember all the original independent credences my peers present me with, so that I can re-input them each time I run the EWV[53]. Unfortunately, this requirement introduces new problems.

3.3.1 Memory Problems

Regardless of how many peers I encounter, to apply the EWV in parallel I must track all their initial credences across time. It is plain to see that this could become a difficult task. My memory for such details is particularly poor, but I think that most people (even those with superior memories), over a long stretch of time or with enough credences to keep track of, would find themselves making a lot of small mistakes. I might forget exactly what *E*'s credence was, or even that *E* had a belief about *P*. I think that after actively utilizing the EWV for a while we might become better at remembering this sort of thing - in the same way that when people actually had to use them it wasn't too difficult to remember each other's phone numbers. All the same, at some point we are certain to hit a wall where we start to forget what everyone's credence in *P* was. If we use the wrong credence, or forget to include it entirely, the output won't reflect what our peers think.

[53] *Ibid*.

Forgetting evidence is not a problem unique to the EWV, but I will note that the type of evidence that is crucial in the EWV, PE, is a form of evidence that is, I think, more readily forgotten than NE, which by its nature is conceptually more tangible.[54] This problem occurs if there is too large a gap between when we encounter beliefs and when we use the EWV. A failing memory can also be a problem if the time gap between using the EWV and reflecting on our beliefs is too large.

Even if we do use parallel processing without making a mistake, if enough time passes, we could forget how we arrived at those (appropriate) beliefs. The details of whose beliefs we have factored in and what those beliefs were would be something we could no longer remember. So, it seems likely that if we attempt to apply the EWV over a long period of time we will often end up with beliefs that are the product of what was a careful but long-forgotten application of the EWV. I have many beliefs derived from analyzing NE that I can't entirely remember my reasons for holding. Usually, however, if I reflect on those beliefs enough, I can retread the reasoning that led me to my belief (or change it). However, my PE about what other people believe feels like something that is more likely to be completely erased from my memory.

This would be another problem for Gardiner, who suggests that when we apply the EWV in parallel we recalculate our original credence by "subtract[ing] the influence of previous peer encounters".[55] While we clearly need some way to re-input our original beliefs into the EWV, I think Gardiner's solution of reverse-engineering them from the EWV's output belief is not the best way to do this. Aside from the memory problems we might encounter, the process of deriving our original beliefs is work that we would be better off avoiding. I will discuss an alternative later in section 3.6.

The source of a belief vanishing is not a problem unique to the EWV, but it does mean our beliefs can become fairly turbulent. Say *A* holds a credence of 0.4 that *P.* She arrived at that credence after applying the EWV to disagreements with *B* and *C,* but now finds that she can't remember what *B* and *C*'s original credences were. In this scenario *A* runs into a problem as

[54] For example, (and this might be idiosyncratic) I find it much easier to remember an argument (NE) than who gave the argument or how convincing they thought it was (PE).

[55] (Gardiner, 2014, p. 12)

soon as she encounters another peer: *D*. Say *D* has a credence of 0.1. *A* wants to assign equal weight to four original credences through the parallel process, but she can't remember what some of those credences were. She could accept that she can't remember everyone's positions and only use those she does remember in her application of the EWV. Given my goal of finding a practical solution to the PPD, this isn't a terrible response at all. *A* is doing the best she can with the evidence she has, and the process will properly recommend a credence based on that information. But perhaps there is a better move for *A* to make.

Adapting the weight of A's credence based on the peers she has encountered is a craftier solution suggested by Gardiner.[56] If *A* can't remember what *B* and *C*'s original credences were, but she can remember that she disagreed with them, she knows that her credence going into the disagreement with *D* is the product of three peers' credences (A+B+C). Thus, she can give her current credence, which is the product of three original credences, a weight of 3:1 relative to *D*'s. The resulting process looks like: $(0.4 * 0.75) + (0.1 * 0.25)$ or $\frac{(0.4*3)+(0.1*1)}{4}$.

Essentially *A* is performing a series application of the EWV, but she avoids the sequence problem by preserving the relative weight of her and her past peers' credences. If she can track the number of peers she has used the EWV with perfectly, this process will produce the same result as re-inputting everyone's original credence. I will call this the weighted series process.

[56] *Ibid*.

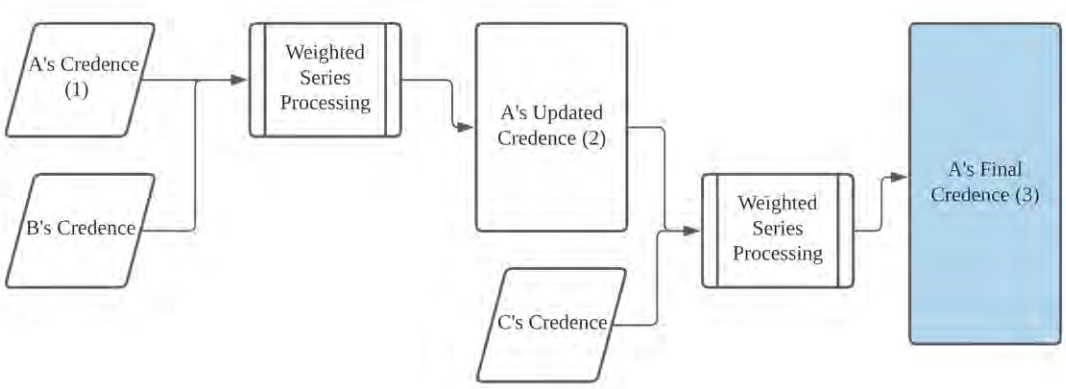

Fig. 5) Weighted Series Processing

The weighted series process is not perfect. Depending on the number of peers involved, the math required by this application of the EWV may be even more complicated than with a parallel application – again requiring a lot of time or a calculator. And, unlike with the series application, we still need to keep track of a variable somewhat external to our belief: the number of peers' credences that we used to arrive at our belief. That is far less of a burden than remembering everyone's credences, but over enough time mistakes are still possible.

For now, it is also difficult to see how the weighted series process could handle disagreement with epistemic superiors or inferiors, something that the parallel process does rather well and a good reason to favor it (I will explain how in section 3.5.4). As Gardiner points out, this process also requires us to know how many peers *our peer* has met.[57] The more peers they have met, the more weight we should afford their views. And if my peer's peers had met with their own peers, that also needs to be factored into the analysis. This regress is not bounded and working around it requires us to have access to information that almost certainly won't be available. The weighted application is also somewhat vulnerable to stubborn peers. If *A* forgets that the stubborn peer's credence has already been counted, she will end up counting it again. This is not as dire as the original problem, but it is still room for what would be to her an untraceable error.

[57] (Elga, 2007, p. 483)

The best way to save a parallel processing application from the problems associated with our limited memories would be for *A* to start keeping track of everyone's credences somewhere. She could do this in something like a peer-belief ledger, so that she never forgets what her peers thought about *P*. In this ledger she would have to keep a record of each peer she encounters and what their credence that *P* is. This is the solution which I favor for the IEWV, and I will discuss this ledger in detail in section 3.5.

3.3.2 Calculation Problems

We should also consider the cognitive overhead that comes with a parallel application. The math involved in assigning equal weight to beliefs to determine the output (recommended) credence is not that complicated. That is, the procedure is relatively easy to understand: simply divide the sum of credences by the number of credences. However, division (and even addition) can become difficult or tedious when enough numbers need to be processed. Thus. the more disagreeing parties involved the more difficult the process becomes. Eventually the process will require so much addition and a feat of division most people will not be able to comfortably perform it without external assistance. We could externalize these processes by using a calculator (or, as I will argue for in Chapter 4, something like a phone app), but requiring an external tool adds to the effective cost of applying the EWV – at least compared to standard applications.

So, if, using the series process, *A* simply assigns equal weight to her and *D*'s *current* credences the process is easy, but she is no longer equally weighting the positions of all of the epistemic peers who came before *D*. She could instead attempt parallel processing, but that quickly becomes overwhelmingly complicated or requires her to remember an unreasonable amount of information. Solving the sequence problem is going to require sacrificing some of the apparent simplicity of the Equal Weight View, which is not ideal given that my goals require the solution to be something humans can use fairly easily. The issue here is that the mental cost of the parallel application can be surprisingly high, and it keeps growing the longer we have been working out what to think about *P* by using the EWV. The weighted series process solves some of these problems, but not all of them, and is less versatile (at dealing with epistemic superiors and inferiors) than the parallel process. Thus, I am going to build the IEWV around the parallel process. I will now move on to

explaining how the IEWV works, and in doing so show how it addressed the election magnetism and skepticism problems.

3.4 Features of the IEWV

IEWV is different from the standard application of the EWV in three areas:

1. The application makes extensive use of a *ledger*, wherein the information required by the application is tracked.
2. It is designed to be applied to cases of disagreement beyond those found in the problem of *peer* disagreement. That is, it can also be used for disagreements with epistemic superiors and inferiors.
3. It has a bifurcated system for handling the processes involved in assessing NE and PE. The credences produced by both processes are retained and used in different circumstances.

The ledger is at the heart of the IEWV. I think that it is easiest to see how this application works by framing how its various components relate to the ledger.

3.5 The Structure of the Ledger

The ledger I describe here is just a metaphor for an organized collection of information. Essentially it is a credence tracking database. In theory all of the information in the ledger could be recorded in our memory. For some especially gifted individuals this is possible, but for most people, due to the amount of information that needs to be recorded, an external ledger will be required for a reliable application. The best way to achieve this would be with the assistance of a custom computer program or spreadsheet (see section 4.2) but to keep the explanation simple I will describe the ledger as a physical book, with a page for each belief that is tested by disagreement.

As I argued in section 3.3, to sustainably perform parallel process applications of the EWV we are required to track the credences of our peers across time, so we don't forget the information needed by the process. Tracking these credences is the primary function of the ledger. With this goal in mind, it is fairly straightforward to determine what the basic layout or structure of the ledger's pages should look like. Each page will have a title, corresponding

to the proposition that is disagreed upon. So, if we disagree about the reality of anthropogenic climate change, the disagreed upon proposition could be stated as [climate change is real], and so the title of my ledger page will be *Climate Change is Real*. Next, I need to know *who* is represented in the ledger so I will add a column to record each person's name. I will refer to the people who have been recorded in the ledger as *interlocutors*, a term that includes epistemic peers, inferiors, and superiors. I also need a column to record each interlocutor's credence that [climate change is real]. To keep everything neat I will also add a column to record the interlocutor's number in the ledger - this isn't necessary to use the ledger, but it will make clear how many interlocutor's credences have been recorded and could make maintaining the ledger a bit easier.

		Climate Change is Real
#	Name	Credence
1	Me	0.9
2	B	0.3
3	C	0.8

Fig. 6) Basic Ledger Page

This ledger is meant to be used to track the positions of epistemic peers, superiors, and inferiors. Thus, it also needs to record the relative epistemic position of each interlocutor. I need to know who is a peer, a superior, and an inferior. Furthermore, not all epistemic superiors or inferiors are my superior or inferior to the same degree, and this is something the ledger can and should be sensitive to. The more someone is my superior (i.e., the more likely I think they are to be correct about *P* than I am) the more their credence ought to be weighted. The inverse applies to inferiors. I will add another column to write down the relative epistemic weight of each credence, which should correspond with the relative epistemic position of the interlocutor it belongs to. I will call the property recorded in this column *expertise*. Expertise is a multiplier, so my superiors will receive a value greater than 1, and my inferiors a value less than 1. I will discuss expertise in detail in section 3.5.4 below. Along with expertise, we want to know when each assessment added to the ledger was made. This just requires adding a date column. The utility of the date column is explained and discussed below.

				Climate Change is Real
#	Name	Credence	Expertise	Date
1	Me	0.9	1	2020.06.14
2	B	0.3	0.3	2020.07.28
3	C	0.8	4	2020.08.15

Fig. 7) Complete Ledger Page

So, each entry on the ledger is made up of five points of data. The name, number and date columns would have us record exactly what each column name suggests; the interlocutor's name, their number in the ledger, and the date that their assessment of *P was made.* The two remaining columns, credence and expertise, are where, aside from the act of physically recording all of this information, the IEWV diverges most sharply from standard approaches. I will now describe in more detail how each of these columns works, why they are necessary, and consider some potential problems. As I mentioned above, the number column is not necessary for the application. However, we do need to know how many interlocutors are in the ledger, so we might as well record that information alongside the entries.

3.5.1 Credences

The credences that are recorded in the ledger are incredibly important for successfully applying the IEWV. They are the raw materials that will be processed to produce the ultimate product of the application - the credence it recommends we hold.

If I take it that someone is my epistemic peer because they are just as likely to be right about *P* as me, then I should assign equal weight to our credences in *P*. But there is some nuance to what I could mean when I say, "these two people, *A* and *B*, are equally likely to be right about *P*".

The most straightforward approach is to say that whatever credences *A* and *B* have about *P* should be weighted equally. The methods that *A* and *B* used to arrive at those credences are not a factor in the analysis. But, as we have seen, if we allow credences that are the product of any method to be used as inputs in the EWV, then the EWV can end up using credences that are themselves *the output of the EWV* as inputs. I think that a lot of the problems that the standard application of the EWV runs into are the result of being too open in its selection of input credences. Specifically, it accepts as inputs credences that are the product of PE analysis (i.e., the EWV).

I will call the credences derived purely from NE analysis "personal assessments". They are personal in that they are formed without giving any consideration to how *other people* have assessed the evidence (i.e., without considering PE). Instead, a personal assessment is formed by evaluating the NE available to you *on your own*. I will call the credences that applications of the EWV recommend, and thus which factor in PE, "electoral assessments".

Allowing electoral assessments to be used as inputs, whether they belong to myself or to others, creates practical problems. The election magnetism problem, discussed in section 2.6.2, is one of these problems. The election magnetism problem only occurs because the population is using the election results (i.e., electoral assessments) to decide who to vote for in their election (i.e., what their input credences are). If only personal assessments are used as inputs for the application, then the 'election results' today can't have a direct impact on which belief wins tomorrow's election. Each election happens independently of the last. The input/output circle is broken, and the electoral assessment loses its magnetism, freeing us from guaranteed epistemic stagnation. Our beliefs can't become untethered from NE evidence analysis if we still have to use NE evidence analysis in the process used to derive those beliefs.

Furthermore, using my interlocutor's electoral assessments also comes with the data collection difficulties discussed in section 3.3.1. We saw that it is difficult to determine whose beliefs served as inputs in someone else's application of the EWV. If my peer's credence was arrived at by applying the EWV (and thus processing PE) then the extent to which their belief threatens my own (via the principle of defeat) is obfuscated. In order to assign the correct weight to a collective assessment, I must know how good my peer is at assessing the epistemic positions of others, how good he is at actually applying the EWV, and who is in his ledger and how they were weighted, so that I don't double count them.[58] Thus, to reliably make use of an electoral assessment as an input to the EWV, I must reverse engineer it to the point that it is just a collection of personal assessments waiting to be processed. If I have this information, then my ledger is functionally just a collection of personal assessments. If I don't have this information then whenever I include an electoral assessment as a credence in my ledger, I am creating blind spots in my dataset through which

[58] In most cases it is unlikely that I will, without some effort, come across this information.

bad or duplicate assessments could corrupt the entire application, and without any mechanism to identify this corruption.

I can avoid these problems by limiting the credences that can be used as inputs for the IEWV to those derived only from an analysis of NE. So, only personal assessments should be used as inputs in the EWV process to determine an electoral assessment output.

This means that the ability to determine what someone's personal assessment about *P* is becomes very important. With some interlocutors this could be a difficult task. Paradoxically, it is made even more difficult the more interlocutors are themselves using some version of the EWV, and thus use electoral assessments. In section 3.6 I suggest how adherents of the IEWV can avoid adding to this problem. In any case, I will still need to be able to differentiate between personal and electoral assessments when my interlocutor does not also follow the IEWV. The extent to which I can do this will surely depend on me asking my interlocutors different questions and my being sensitive to different cues in how they express their positions.[59] I could also guide my interlocutor into making a personal assessment in our conversation (in a way similar to how a professor might guide her student into making his own assessment). I think that being able to differentiate between these two types of assessments will have to involve many subtle processes and can't be clearly outlined here. I suggest appealing to heuristics as a solution to this problem in section 4.3.

3.5.2 Names

The IEWV is focused on collecting personal assessments of NE. These assessments are usually produced by the minds of individual thinkers, thus the source of credences in the ledger ought to be individuals. However, we often hear about assessments that are supposedly made by institutions or groups (e.g. the World Health Organization recommends doing *x*, some newspaper endorses candidate *y*, a group of scientists advocates for position *z*). I will refer to these as *collective assessments*.

At first glance, there is a lot of utility to be gained by factoring collective assessments into the ledger. Collective assessments represent many people, so I don't have to track down and

[59] Different, that is, from how we normally engage in conversations about beliefs.

record lots of different sources to gain something of the same epistemic weight. Their expertise value might be easier to assess than that of individuals. As a layperson I suspect I have a better chance of working out how trustworthy a specific science publication or advocacy group is than the trustworthiness of 30 individual scientists. It is easier to find discussion about and endorsements of groups and organizations than individuals.

However, collective assessments come with some problems. It is not always clear what the relationships between the collective assessment and its constitutive personal assessments are. If I see that a group of scientists has signed an open letter, then perhaps I can assume that each signee probably agrees with the contents of that letter (or at least most of the letter). That is, the letter reflects their personal assessments. But if an institution states that their position is that *P*, it might be the case that it is only the view of the higher-ups in that institution, or even that the view is being presented for reasons that aren't at all epistemic (e.g., for political or economic reasons).

Furthermore, if I don't know *who* the collective assessment has been made by, I may end up counting some sources more than once in my ledger, thus giving them more epistemic weight than they are owed. For example, recoding the scientist Pete's personal assessment in one entry, and unwittingly recording Pete's department's assessment (which, of course, includes Pete's assessment) in another. This is similar to the problem we encounter when using electoral assessments as input credences just outlined in section 3.5.1. In the same vein, assigning an actual expertise value to a collective assessment could also be very difficult to do. This is the case if it is difficult to determine how the assessment was made, who contributed to it, what their individual expertise levels are, or what motivations the collective and its members might have. It might be easy enough to work out that the source of a collective assessment is trustworthy, but without details about its constituents working out how to weigh the assessment relative to others is tricky.

Sometimes collective assessments are simply a belief that is shared by all of its constituent members, but there might also be disagreements among those members. From the standpoint of someone applying the IEWV, the best-case scenario is that such collective assessment was also formed using the IEWV. But, as I just argued, I want to avoid relying on other people's

applications of the EWV because doing so can quickly result in an opaque data set that is difficult to work with without encountering problems.

Given these issues, rather than using collective assessments as unique entries in the ledger, I would be better off treating them as bibliographies; as tools that point to various personal assessments that I could include in my ledger.

These problems don't mean that collective assessments are useless. With some skill, effort, and the right methodology, we could probably determine the expertise value of a collective assessment well enough to make use of them. However, because assessing a collective assessment is usually going to require a different kind of research to assessing a personal assessment, and is ultimately pointless if we could instead record its personal assessments directly, I will not factor them into the IEWV. Of course, outside of the idealized resource-insensitive circumstances assumed for this application, collective assessments become increasingly attractive, and I will argue that they are a very valuable tool to reduce the amount of time needed to perform the EWV in Chapter 4.

3.5.3 Dates

The date column is for the date that the interlocutor *made their assessment*, not the date that I first encountered or recorded it. For example, if I watched a documentary from the 1980s about climate change *yesterday and* am *today* recording the views of a scientist in that documentary to my ledger, I should record his views as coming from the *1980s*, not today (when I made the recording), or yesterday (when I encountered his views). Ideally, the date of assessment also isn't the date that the source *communicates* their credence. If I express today that I think free will is impossible but haven't really thought about the problem of free will since my undergraduate classes, then the assessment ought to be dated as coming from when I was in those classes, not today. We want to know when the assessment was made because that will impact the expertise value of the assessment over time (see section 3.5.4 below for details).

Basically, if an assessment becomes 'outdated' - because *new* evidence has arisen that was not part of that assessment - it becomes less likely that it is the best assessment, and our ledger ought to reflect this. Note that assessments become outdated because they weren't

factoring all the evidence available to us now, not because they are old. For questions where no new evidence has become available since ancient times, the *age* of ancient assessments does not count against them.

Seeing that a view is old may also prompt me to seek out an updated assessment from the same interlocutor. We probably aren't ever going to have perfect knowledge of when an assessment was made. I can't remember when a lot of my own beliefs were formed, much less those of others. However, it is still a data point that we should do *our best* to track (again, it is okay that we can't be perfect here, we just need to do the best we can to get the most utility out of the process).

3.5.4 Expertise

A strength of the EWV discussed in section 2.5.5 was how easily it is applied to the problem of disagreement in general (i.e., beyond dealing *only* with epistemic peers). The two other relative epistemic positions an interlocutor can hold are epistemic superiority and epistemic inferiority. It is quite intuitive what the EWV would prescribe for me to do when I encounter epistemic superiors and inferiors with views about *P*. I should assign more epistemic weight to the credences of my epistemic superiors than to my own, while assigning less to those of my epistemic inferiors.[60] I should adjust the weight ratios based on just how much more they are my epistemic inferiors or superiors. The *expertise* value assigned to each interlocutor in the ledger is how I assign this weight. I will use myself as the base, with an expertise value of 1.[61] Of course, my epistemic peers will also have their credences weighted at 1. My epistemic superiors will have their credences weighted greater than 1, and my inferiors less than 1. How much that value is greater or less than 1 is determined in proportion to the degree to which they are my superior or inferior respectively.

So far this looks like a nice system where we, in line with intuition, end up weighting the beliefs of experts (i.e., the people most likely to be right about *P*) above those of lay-people.

---

[60] This is mentioned by Carey and Matheson (2013, p. 10) where they state, of their epistemic election, that "the relevant experts' opinions count for more". Also in (Elga, 2007).

[61] We could just as easily use a scale using someone else as its base (e.g., a layperson or expert), but I think keeping ourselves at the center of our process makes the idea a little easier to grasp. Using a different base is an important step in developing the PEWV in Chapter 4.

Thus, the electoral assessment produced by this process will lean towards reflecting the beliefs of experts more than those of lay folk or the entirely uninformed.

However, as Elga mentions, it can be difficult to work out how much of an expert someone really is.[62] Determining the appropriate expertise value of an interlocutor is often going to be a difficult task, especially with superiors. It is hard enough to determine whether someone is an epistemic peer, superior, or inferior. The task becomes harder if I also need to determine the *extent* to which my interlocutor is my superior or inferior. And this is just what the expertise column asks me to do.

The problem is not too great when I want to work out how much someone is my epistemic inferior (provided they are in fact my inferior). Usually when you do in fact know better than your interlocutors it is easy enough to see. For example, a professor should be able to spot where her student has made a mistake or lacks knowledge or understanding fairly easily. But why should we care about the views of our epistemic inferiors at all – wouldn't counting them just make us more likely to be wrong? I think that what inferiors think can be valuable PE, as long as we accurately assign an expertise value to their credence. If enough of my inferiors disagree with me, then their combined epistemic weight should cause me to revise my belief.[63] Even if I think I am twice as likely as my 10 inferiors to be right about *P*, that only gives my credence, relative to theirs, a weight ratio of 2:10. I think it is fairly intuitive that I would have a reason to revise my belief if all 10 of those inferiors disagreed with me.

In domains where I am a novice, determining the expertise of people more likely to be correct than myself can be awfully tricky. I think that most people make mistakes here, and sometimes we make really big mistakes. This could explain why we occasionally see people buying into the positions espoused by fringe pseudo-experts (i.e., people who are not actually very likely to be correct) despite their views not being held by the mainstream of their field. In many cases these "experts" are even considered outcasts or fools by the rest of their field. Yet, the way in which the pseudo-experts present themselves is such that a layperson cannot easily tell them apart from real experts. Recall that we also want radical advancements to be

[62] (Elga, 2007, p. 483)

[63] This is most clearly seen when they are only barely my inferiors.

possible. Experts can and do mistake revolutionary advancement for quackery, so it is unsurprising that we, as laypeople, often make the inverse mistake. There is also a difference between an actual expert with a radical theory that happens to be wrong and a pseudo-expert. The expert is likely to be correct but happens to have a false belief – the pseudo-expert is unlikely to be correct and unsurprisingly has a false belief.

While assessing expertise is difficult, I do wonder if it is more difficult to assess expertise than it is to get to the right answer relying only on our own analysis of lower order evidence. For questions that require expertise that we lack we have no choice but to find an expert to rely on. Thus, it might very well be that the answer depends on the difficulty of the problem we are dealing with. That answer is not something we have much access to, because, as is discussed by Morton, assessing the difficulty of any problem is fraught with its own difficulties.[64]

If we buy into the principles of the EWV we should not be too dissuaded by the problems associated with expertise. In fact, the importance of assessments of expertise in the IEWV is a reason to favor it. An adherent of the IEWV must come to value different intellectual capacities than those valued in NE-assessment-based belief generation models. If all of my beliefs are generated by my interpretation of the NE, then my ability to assess other people's interpretation of the evidence is not that important. In fact, it has essentially no bearing on whether I will be right or wrong about the rest of my beliefs (because what other people think is not a factor in my thinking). This is the case in the SV solution to the PPD. But, when I buy into the EWV, being able to assess relative epistemic position accurately becomes incredibly important. It is a big part of my PE. Thus, it becomes crucial to be a good judge of the epistemic standing of your interlocutors.

The focus on assessing interlocutors does not strike me as particularly unusual when compared to contemporary epistemic practices. In the age of fake news and online propaganda, being able to accurately assess the quality of information sources has become incredibly important. Because so many of our beliefs are based on the mediated testimony we find in books, on the radio, TV, or internet, we need to be able to tell apart trustworthy and

[64] (Morton, 2012, p. 88)

untrustworthy sources. This is also true when it comes to making assessments in all sorts of fields where we do not take ourselves to be anywhere near experts and have no hope or intention of ever acquiring expertise. For example, I have no plans on educating myself on nutritional science enough to consider my own thoughts in this area worth mentioning next to those of actual experts, but I do want to optimize my diet. Rather than become a dietician or nutritionist, it is more efficient for me to spend (much less) time learning about what the current field of nutrition science looks like, and thus who in that field can be trusted to give me good advice about what I should eat. Effective use of the Equal Weight View simply requires that we focus some of our attention on assessing the expertise of *all* the interlocutors whom we want to use in our epistemic election. I will discuss ways to do this in Chapter 4.

Unfortunately, there are also mechanical problems that come with using expertise in the application. The expertise scale I suggest uses the person keeping the ledger as its base - giving them an expertise value of 1. This is useful because it means I only need to think about how much more or less likely than I am my interlocutor is to be right. However, if my level of expertise changes (say I learn more about the question, or realize I knew less than I thought) then I either need to change the expertise of everyone else in the ledger relative to myself or change my own expertise to more or less than 1. Changing my own expertise value is easy to do but could make it harder to assess the expertise of new interlocutors in proportion to the old. Alternatively, recalculating all of my interlocutors' expertise will make new assessments a relatively straightforward procedure, but will be a lot of work upfront. Neither solution is ideal, but both get the job done. I think that doing the work upfront by recalculating or reassessing everyone's expertise is probably the better move - it is a good idea to do some maintenance on the ledger periodically anyway, especially for expertise values which are difficult to nail down.

The addition of the expertise variable also exacerbates the familiar problem of the EWV being too difficult to calculate. Now we are multiplying each person's credence by their expertise and then aggregating those values, so it quickly becomes too complicated for people to easily do.

Despite the difficulty associated with tracking expertise, I think that something like it is necessary if we want to use this theory in the real world. Even if cases of peer disagreement

weren't so rare, they will almost never be *all* of the disagreements I have about *P*. Resolving *peer disagreement* isn't achieving much if it means ignoring all of the PE associated with other interlocutors on *P*, especially when I think many of them are my (and my peers') epistemic superiors. The resulting belief would always be based on less than all of my evidence, and thus isn't very attractive. To apply the EWV to real-world disagreement and get the most utility from it we must factor in our interlocutors' expertise. Again, I will consider some easier ways to do this in Chapter 4.

3.6 Idealized Methodology

In this section I will explain the methodology of the IEWV application. That is, how it moves from evidence to beliefs. When discussing credences in section 3.5.1 I mentioned that in the ideal application of the EWV I need to track both my *personal assessment* of the NE and the *electoral assessment* produced by using the EWV (by processing PE). This means that when I use the EWV I will not *replace* my original belief (i.e., my personal assessment) with the output of the EWV (i.e., the electoral assessment). Thus, I should know what both my personal and electoral assessments are. Below are three flowcharts that should help explain what this change means for the IEWV.

Fig. 8) Personal Processing

A basic process of belief formation is illustrated in figure 8. A NE input goes to a process of NE analysis which produces a belief output. In essence this is the belief formation methodology the SV advocates for. I'll call this 'personal processing'. The details of how the NE analysis is conducted are not important for my purposes here. It simply involves whatever method of information processing is used on the NE – the possibilities range from logical deduction to heuristic pattern recognition. Ideally the process of NE analysis is something that accurately and reliably determines *what the evidence points to*. But, in the context of a discussion about the EWV, all that really matters is that it doesn't involve any PE. Naturally

the belief produced by personal processing is also a personal assessment. As we move into the more complicated EWV methodologies, it will be useful to think of personal processing as a module in our belief-forming system.

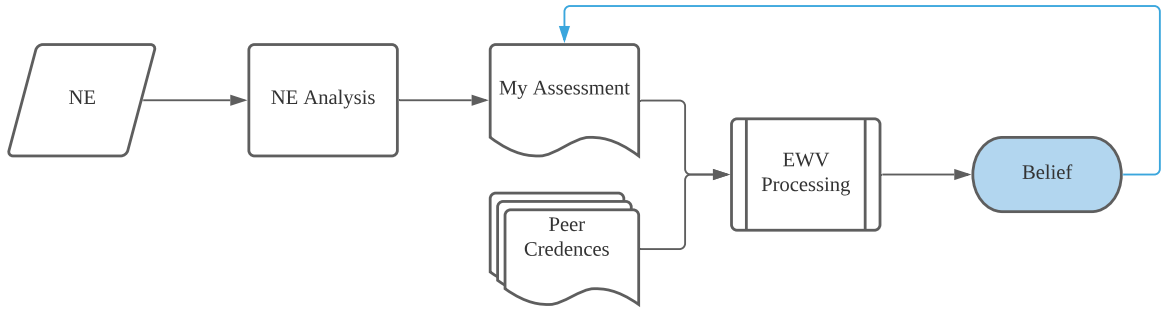

Fig. 9) Equal Weight View Processing

In figure 9 we see the belief formation process utilized in standard applications of the EWV, assuming they are also using the parallel PE processing outlined in section 3.2. The process begins with the personal processing module just discussed. The NE analysis produces a personal assessment which, along with my peers' credences, is processed by the EWV. This is done in parallel, so whenever I 'run' the EWV I simply add all of the credences together, then divide by the number of credences I compiled.

This flowchart highlights some of the problems with the EWV I have already discussed. Note that the output of NE analysis is labeled as 'my assessment' rather than 'personal assessment'. This is because this credence is *replaced* by the output belief of the entire process after it has been run. That belief is the product of the EWV, and so is an electoral assessment. Because I replace my original personal assessment with this electoral assessment, I am still vulnerable to the sequence and stubborn peer problems, even though I am using a parallel EWV process. Thus, the method does not reflect the principle of equal weight. While it uses a parallel process, the method assumes that I will simply remember all of my peers' credences and be able to calculate the average without assistance. This makes it quite likely that I will make a mistake, either by misremembering or miscalculating. Also, note that the credences being processed are only those of my epistemic peers. The method does not track expertise or an equivalent property, and so has limited applicability.

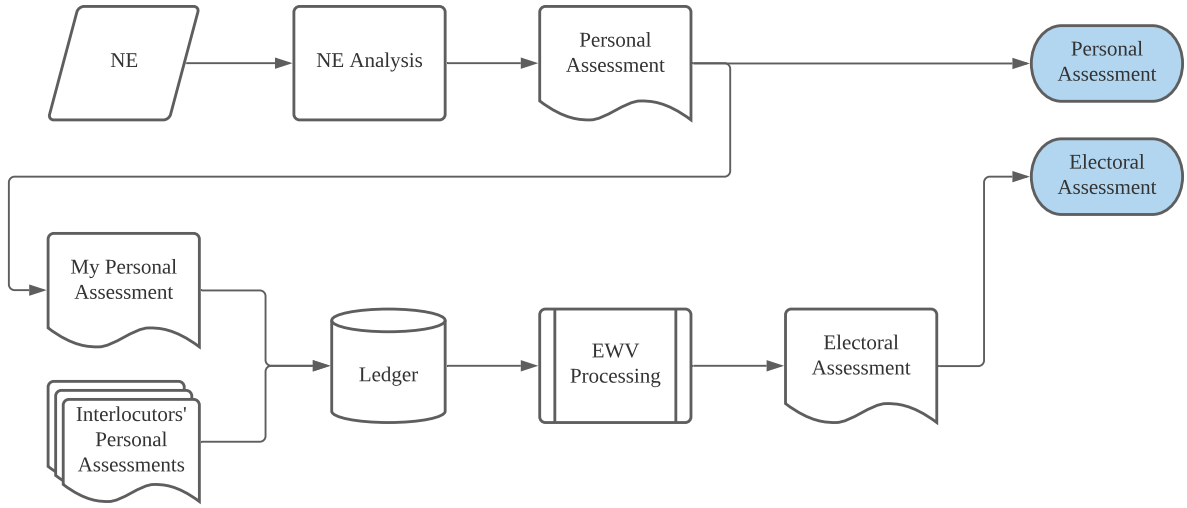

Fig. 10) Idealized Equal Weight View Processing

In figure 10 we see the processes involved in an application of the idealized EWV. The initial analysis of NE is the same in the SV and produces a personal assessment output. Note that, unlike in the EWV, this assessment, *my assessment*, is not replaced by the output of the EWV. The only path to *my* assessment is through NE analysis. My personal assessment, and those of my interlocutors, are recorded in the ledger. The ledger acts as a database and includes the information described above (names, credences, expertise, and dates). Tracking expertise allows the process to accept the credences of all interlocutors, rather than only my epistemic peers. So, the process also incorporates the views of epistemic superiors and inferiors.

The ledger is the input for the EWV process. First, the process multiplies each credence with its expertise value to produce a weighted credence. Then the sum of the weighted credences is divided by the sum of the expertise values in the ledger.[65] The result is my electoral assessment, which has no effect on my personal assessment. So, my electoral assessment equals $\frac{(\mathrm{A}*\mathrm{Ae})+(\mathrm{B}*\mathrm{Be})+(\mathrm{C}*\mathrm{Ce})}{(Sum\ of\ e)}$ (where 'e' is expertise, 'Ae' is A's expertise).

[65] Because the credence weights have been multiplied by expertise values, the process needs to divide by the sum of expertise values instead of the number of interlocutors.

Notice that both my personal assessment and my electoral assessment are labeled as outputs. For the purposes of this thesis, I am not particularly interested in the metaphysical nature of personal and electoral assessments. That is, the question of which assessment is what I 'believe', or if I can 'believe' both assessments. I am concerned with mental states that will inform my actions, whether those actions are directed outward into manipulating my body to achieve something in the world, or inward towards reshaping the contents of my mind. Recall Hob's friend Fae. By this model, Fae's personal assessment might be that global warming isn't real while her electoral assessment is that global warming is real. The question of what Fae truly "believes" in this situation is philosophically interesting, but I don't think it affects Fae's ability to *use* both assessments where they are appropriate. So, where should Fae use her two assessments?

Assuming the principles underlying the EWV, the electoral assessment produced by the idealized application is more likely to be correct than Fae's personal assessment (that is because it is better supported by *all* of the evidence available to Fae (both PE and NE)). Thus, Fae ought to use the electoral assessment to guide her actions. In this case that means Fae should *act* as if [global warming is real] is true (which is her electoral assessment). By doing so she is acting on a 'belief' that is, from her perspective, her best 'map of reality'.

However, there are still areas where Fae should use her personal assessment. I have argued that the EWV process should only accept personal assessments as inputs. In real terms this means that Fae should maintain and present her personal assessment when she is engaging in interpersonal epistemic activity (e.g., discussing global warming with her friend). That is, when we are engaged in the epistemic activity of working out *what to think* we should present our personal assessments of the questions at hand.[66] Thus, it is also going to be important for Fae to revisit and update her personal assessments, to account for new evidence she acquires and developments in her capacity to assess that evidence.

[66] I think that tis bifurcation is in line with Feldman (2007, p. 17) when he says: "This skeptical conclusion does not imply that people should stop defending the views that seem right to them. It may be that the search for the truth is most successful if people argue for the things that seem true to them. But one can do that without being epistemically justified in believing that one's view is correct."

In this way the idealized EWV looks a lot like the SV. The independent thinking required for epistemic progress is still strongly preserved in this model. Rather than dissolve into the electoral hive mind, I become an active agent in my epistemic community. Part of my responsibility as a member of that community is to continue to produce personal assessments to the best of my ability, and to share those assessments with the community. If everyone does so, the community as a whole, through applying the IEWV, can arrive at its best electoral assessment.

So, if I am a scientist who has made a radical breakthrough that I am convinced of, but none of my peers agree with me, there is no reason for me to abandon my breakthrough (as there may be with the standard EWV). Instead, I keep doing my work and trying to convince my peers to change their personal assessments of the NE. However, when it comes time to *act,* I should act as if my electoral assessment were true.

Thus, a nation of IEWV adherents perfectly aware of each other would all act as if *P* were true if that were their shared electoral assessment,[67] because they think that it is their best assessment. They could possess very different personal assessments and should openly argue about them. Nonetheless, they would act in accordance with their electoral assessment, until that electoral assessment changed. This behavior might seem very strange to us, but I suspect deferring to the collective in this way isn't so unusual in some parts of the world – though that is a tangent I won't venture into. In a sense the IEWV turns us into part of a collective mind. Each individual tugs the mind in the direction of their personal assessment by being an active epistemic citizen, but ultimately recognize that the electoral assessment that they are building together is more likely to be true,[68] and thus ought to guide their *actions*.

3.7 Conclusions

To wrap up the discussion of the IEWV, I will reflect on how well the view has achieved its goals. Remember, the IEWV is not intended to serve as a *practical* application of the EWV. Rather, it is meant to be a version of the EWV that *works* when applied to the kind of

[67] By "perfectly aware" I mean that they all know exactly what each other's credences and true expertise values are.

[68] Based on the principles of the EWV they buy into.

disagreements that we see in the real world, even if it is too cumbersome. If it works, we can use it as a schematic for devising the PEWV in Chapter 4. Often real-world disagreements involve encountering many people who disagree with us (not all of them our peers) over an extended period of time. The standard EWV's inability to handle such disagreement meant that it was poorly suited for its *actual use case* and so had serious 'external practical problems'. How successful is the IEWV at overcoming these problems?

3.7.1 Overcoming External Practical Problems

We saw that the IEWV does not encounter the sequence problem because, instead of processing credences in series, it uses a parallel process. Because 'running' the IEWV does not affect any of its credence inputs (including the users) the order interlocutors are encountered in does not matter. Adopting a parallel process also protects the view from the problem of the stubborn peer by preventing that peer from mechanically occupying more than one peer position.

Next, we have the election magnetism and skepticism problems. The election magnetism problem was the result of EWV users systematically replacing their personal assessments with the electoral assessments produced by the EWV. The IEWV, by bifurcating 'beliefs' into personal and electoral assessments, has insulated personal assessments from the electoral assessments. Only beliefs generated by analysis of NE are used as inputs by the IEWV. Thus, the 'winner' of the epistemic election isn't going to have a magnetic field around it. The IEWV acts as an aggregator of the epistemic communities thinking, rather than a replacement for that thinking. The bifurcation of beliefs also provides some protection from the threat of skepticism (which we saw was really just agnosticism about some beliefs). The electoral assessment produced by the IEWV may indeed be agnostic about *P*, but I am also meant to maintain and use my personal assessment in some scenarios, and that assessment is only going to be as inclined towards agnosticism as I am. Again, I am not going to explore which of these assessments is what I 'truly believe'. If my personal assessment is my true belief, then the IEWV just recommends that I *act* as if I were more agnostic sometimes.[69] If my

[69] Because, while I may not believe it, the totality of my evidence doesn't point towards certainty.

electoral assessment is my true belief, then the IEWV just asks that when I am engaged in certain epistemic activities, I use my personal assessments instead of what I believe.[70]

So, the IEWV, by adopting a parallel processing model that only uses personal assessments as inputs, has overcome the major external practical problems facing the EWV.

By including the expertise variable, the IEWV opens itself up to processing disagreements with all people, not just among epistemic peers. This is attractive because it inclines us to side with our epistemic superiors.

3.7.2 New Internal Practical Problems

While the IEWV avoids many of the problems of standard EWV applications, it does not do so without paying a price. The IEWV was designed to be useable by humans, and I don't think it asks the impossible, but 'running' the application is not a cognitively cheap procedure.

In order to solve the memory problem that comes with using a parallel credence processing method, the IEWV was forced to create a disagreement ledger. Unfortunately, the ledger makes the view somewhat like the tedious marathon training plan. The first problem is that the ledger is something literally outside of my mind. Consider what happens when I learn that *B* disagrees with me. The automatic process of committing his belief to memory is replaced with the annoying task of readying the ledger, finding or creating the page for the belief in question, and then physically writing down all of the necessary information. Perhaps this doesn't sound too demanding, but I suspect that I would only be willing to go through this effort for a very tiny set of my beliefs. The mere barrier of having to input information into an external system has already greatly reduced the utility of the IEWV by decimating the number of beliefs we could realistically expect it to be used for.

[70] By this model my personal assessments occupy a middle ground in my belief generation process. The IEWV treats epistemic engagement with others as an act of collective belief generation which starts from that middle ground.

There is also a worry about how reliably we can collect the information required by the ledger. The first new difficulty with the IEWV is that I am only meant to record personal assessments in it. That means I need a method for distinguishing my interlocutor's personal assessments from his electoral assessments. I also need a way to assess my interlocutors' level of expertise, which might be a difficult task. Ideally, I will also know when they made their assessment. These data points all make the IEWV less user-friendly. First, the process of collecting this information, whatever it is, is more (potentially very difficult) cognitive labor. Second, writing down all of this information takes time and energy.

Finally, the IEWV is much more complicated to perform than the standard EWV. Processing the information in the ledger is going to be too difficult for most people to do in their heads, so we will have to use calculators. While not necessarily difficult, this process takes more time and effort.

The IEWV might never demand us to lift an impossibly large cognitive load, but it is still a very taxing belief-forming system. There might be some scenarios where the IEWV's utility outweighs the time and effort required to use it. I suspect that those use cases will be most common with large institutions like businesses or governments, where arriving at the best assessment is very important. Individuals would only bother with the IEWV for beliefs they really cared about.

So, in the process of overcoming the EWV's external practical problems, the IEWV has created a number of internal practical problems for itself. While the view does account for the *limits* of our capacities, it is not sensitive to how unwilling most people would be to pay the cognitive costs associated with its application. I will now move on to developing the practical Equal Weight View (PEWV) as a methodology that greatly reduces these costs without undermining the functioning of the IEWV's various modules.

***

Chapter 4. A Practical Equal Weight View

4.1 Introduction

My aim with this thesis has been to work out how we ought to apply the Equal Weight View (EWV) in the real world. As it is usually presented, the EWV fails to resolve scenarios that represent real-world disagreements in accordance with its own principles. In Chapter 3 I developed the Idealized Equal Weight View (IEWV) to resolve the underlying problems with the standard EWV which caused it to violate its own principles when applied to real-world cases of disagreement. The IEWV overcame these problems but did so at the cost of the application becoming incredibly resource-intensive. Unfortunately, the cognitive cost of using the IEWV is so great that it is difficult to recommend using it in day-to-day life. Most people would quickly give up on the application, and those who don't would have to spend a lot of their time using it. Thus, while the IEWV is a step towards practical utility for the EWV, there is still more work to be done if we are to arrive at a real-world application.

In this chapter I hope to finally arrive at an actionable version of the EWV that works for real-world disagreement scenarios. I call this application the practical Equal Weight View (PEWV). My methodology for getting to the PEWV is quite simple: start with the IEWV as a model application and convert it into something less demanding to use. I will attempt this conversion by isolating various parts of the IEWV and attempting to reduce their cognitive cost.

I use IEWV as a starting point for the PEWV, rather than simpler standard applications, because the view needs to work with real-world disagreements. As we saw in Chapter 2 this is something standard applications fail to do. The IEWV does work with real-world disagreements (i.e., it resolves these disagreements without violating its principles), but is difficult to perform. To reuse the terminology developed in section 1.3, while it does not have the standard EWV's ex*ternal practical problems*, it does have serious *internal practical problems.*

A lot of work goes into applying the IEWV. There are many steps involved in using the IEWV, as was shown in figure 10. My strategy in this chapter will be to look for ways to reduce the cognitive cost associated with each step or process of the IEWV, without

compromising its ability to perform its function. If the application becomes efficient enough it could work in the real world and be easily utilized by normal people.

The first part of the IEWV involves whatever processes we use to generate our personal assessments. It is important that we continue generating personal assessments to avoid the election magnetism problem (discussed in section 2.6.2 and solved in 3.2), so this step is unavoidable.[71]

Next, the relevant information about our interlocutors needs to be collected. This includes everything that goes into working out what should be recorded in the IEWV's ledger. That is, *who* is recorded in the ledger, what their *credence* and *expertise* value is, and *when* they made their assessment. The biggest hurdles with data collection are making assessments about an interlocutor's expertise and credence.

After the information is collected it must be stored. This process covers all of the steps involved in actually making a ledger page. This includes setting up a page in the ledger, writing down the relevant data, and updating the ledger when necessary. Of course, the fact that all of this information is stored in an external ledger is an important part of the IEWV. Having to open up the ledger and write down an entry is time-consuming. While it isn't the most difficult task involved in the IEWV, the friction involved in performing daily (if not more frequent) data entry is where most people would throw their hands up and stop using the IEWV.

Finally, the information in the ledger is computed to arrive at the recommended credence. Here we are trying to use our PE, as it is recorded, to get to an electoral assessment. In the IEWV the process involves calculating the sum of the credences multiplied by their expertise values and dividing this number by the sum of the expertise values ($\frac{(\mathrm{A}*\mathrm{Ae})+(\mathrm{B}*\mathrm{Be})+(\mathrm{C}*\mathrm{Ce})}{(Sum\ of\ e)}$). The

[71] Perhaps in cases where, compared to my interlocutors, my expertise is very low there is little utility to be gained from me tracking my personal assessment. For example, being a layperson, what I think about the cause of global warming is barely a factor in my electoral assessment.

math involved isn't very advanced, but it is difficult, and will usually require using a calculator, especially after more than a handful of interlocutors have been met.

There are many possible approaches to reducing the epistemic cost of using each of these modules. In this chapter I am going to explore two of them. The first approach, which I will call *externalization*, leans into the fact that the ledger exists as an object outside of our minds. By utilizing modern computing power, we can greatly reduce the cognitive load associated with storing and processing ledger data. The second approach, which I will call *simplification*, goes in the opposite direction. The simplification approach attempts to change the IEWV so that it can be applied without requiring any external assistance. This approach makes use of heuristics as cost-effective alternatives to the rigid calculations required by the IEWV.

I am not sure that either approach really gets us all the way to a PEWV. While neither application is suitable for resolving every instance of disagreement, they are applicable to a far wider range of real-world disagreements than the standard or Idealized EWV applications are. The approaches also complement each other and can be used in conjunction with one another. Externalization produces an application that has a high fidelity to the IEWV but is still somewhat expensive. Simplification produces (what is likely to be) a less accurate application, but one that is also very cheap. We can default to using the simplified application for day-to-day disagreements while saving the more expensive externalized application for cases we really care about.

4.2 The Externalization Approach

The externalization approach embraces the fact that the ledger exists outside of our minds. The IEWV moves the record of credences from my memories to the ledger and the processing of that ledger to a calculator, but these processes could be more efficient. By making these processes more efficient the application becomes easier to use.

An obvious flaw with the IEWV was that it required us to keep a paper ledger. I think that this physical medium was best for explaining the various mechanics at play in the IEWV in Chapter 3, but it needlessly holds back the application. Given the technology available to us today, we would be much better off making the ledger *digital* and storing it on a computer.

After all, the ledger is essentially a collection of spreadsheets. I think the solution lies in repackaging the spreadsheet in something that is easier to interact with – something like a mobile phone application. So, instead of writing in a book or doing data input on Excel, we should just use an app on our phones. Below I map out what this app, which I will call "Weightify", could look like.

The main reason to use something like Weightify is to reduce the cost of inputting data into the ledger, which is currently the external application's greatest cognitive bottleneck. Just how much of the cognitive load the application is capable of removing depends largely on how it is designed. I don't want to spend too much time talking about these design features, but I think that they are important for understanding how further externalizing the IEWV could work.

The Weightify app has some immediate advantages over the original paper ledger. As digital spreadsheets are relatively small files, they are easy to copy, share, and backup. If backed up properly we would be, if we accept Andy Clark and David Chalmers' arguments, protected against losing a vital part of our extended mind.[72] The spreadsheet is also much easier to edit, and thus ledger maintenance would be a less taxing affair. Indeed, the input and editing processes can become quite refined. Most importantly, when it comes to applying the view, Weightify would allow for a significant amount of automation. In fact, the entire data processing module could be automated with some basic programming. Multiplying credences with expertise values, adding the results, and dividing that by the sum of expertise values can all happen automatically and instantly. This means that one of the most taxing parts of the IEWV, data processing, has essentially no direct cognitive cost in the externalized PEWV.

Because the processing module is external, we need to somehow access its output. The simplest way to do this is by reading the output – I don't think this is much of a burden. Many people already use computers to quickly generate beliefs, whether by using a calculator, a search engine like Google, or even asking a digital assistant like Siri. As long as

[72] (Clark & Chalmers, 1998)

the system is reliable, fast, and is cognitively cheap to use we are happy to retrieve information from, and generate beliefs with, external devices every day.

The externalized PEWV doesn't demand us to perform complicated computations, but it does require us to build and maintain the external ledger. Unfortunately, this means that the storage module is still problematic for the view. While a digital file might be easier to preserve, there remains a barrier between it and our thoughts. That barrier becomes an obstacle when we are trying to add information to the ledger. Opening up a spreadsheet on the app to do some manual data entry every time we meet an interlocutor with a view on *P* is still tedious, and it isn't something we will be willing to do for a lot of beliefs.

In the IEWV each ledger entry is made up of five points of data: an entry number, a name, a credence, an expertise value, and a date that is meant to correspond with when the assessment was made. I think that all versions of the PEWV will eliminate the date data point as a requirement. While this information can be useful for maintaining the ledger, it is ultimately metadata that isn't necessary for EWV processing. Some of the utility of the date recording could be maintained by making it an optional tag (to be used when the date of an assessment stands out in some way), but it shouldn't be required for every entry. The entry number could also be added automatically by Weightify but given the app can instantly count the ledger entries it is now unnecessary. That means that once a page is set up, we only need to input three points of data per entry; a credence, an expertise score, and a name.

A first step towards making data entry easier would be to design the app so that credence and expertise are recorded at the same time. If on a grid the x-axis corresponds with credence and the y-axis corresponds with expertise, then we just need to select a point on the grid to indicate what the credence and expertise of an entry are. Instead of typing two entries, a single tap of a finger is all that would be required.

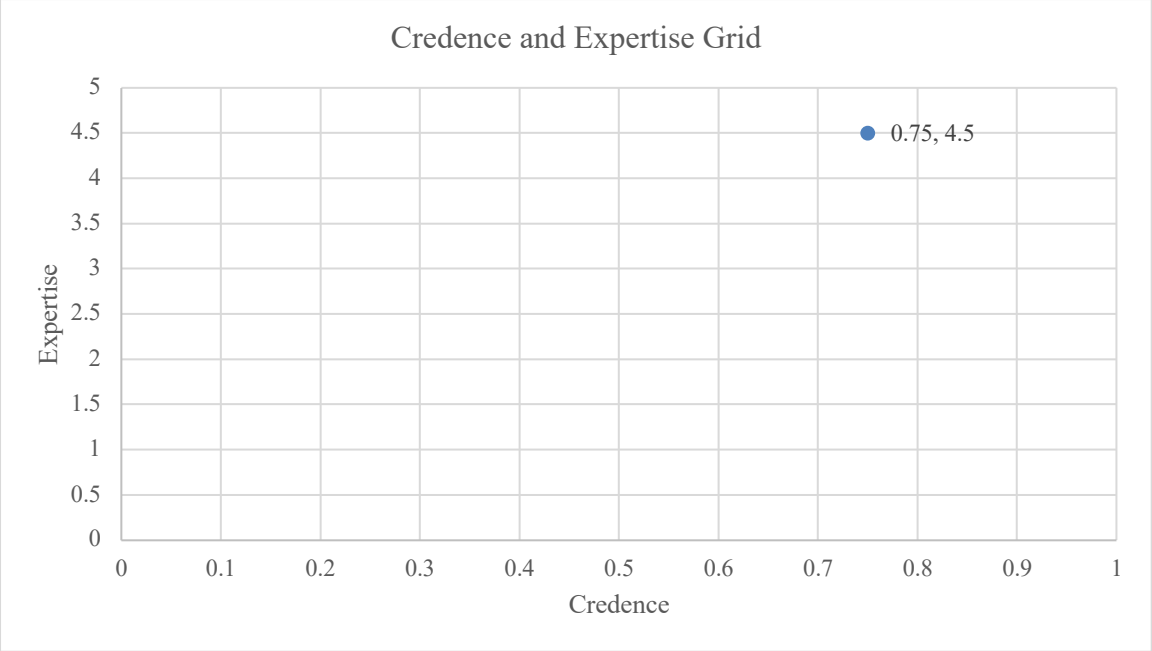

Fig. 11) Credence and Expertise One-Touch Input

In the IEWV both the expertise and credence values are fine-grained. This detail entails a small cost in the storage module; we have to be precise when we input information. Worse, being required to make precise assessments makes collecting information about our interlocutors more difficult. The precise credences of our interlocutors are difficult to ascertain at the best of times. It isn't always clear what the difference between a 0.4 and 0.3 credence is, never mind detail beyond one decimal place. The process of determining expertise, as we saw in section 3.5.4, has even more problems associated with it. Fortunately, using a computer to perform our data processing and storage creates some room to work around these difficulties. If we allow the computer to handle the numbers for us, we greatly simplify our task.

Working out where someone should be on a fine-grained expertise or credence scale is difficult. We can make our task easier by making those grains a bit larger and more tangible. Instead of trying to work out exactly what value our interlocutor's credence is, we can focus on identifying roughly what their attitude towards *P* is. Exactly how we break down these attitudes is something we can be flexible with, using more or less detail depending on our confidence in assessing credences. An example of a set of these attitudes could include 'agnostic', 'weak inclination', 'strong inclination', and 'certainty' that *P* is true or false. That

leaves us with 7 attitudes to worry about (compared to the 11 of a single-digit numeric credence assessment). Unlike fine-grained credence assessments, I think it is quite clear how all of these positions actually differ, and it should be much easier to put an interlocutor into one of these boxes than it is to assign them a credence score. At least, this is the case for new users – it is possible that with enough practice assigning credence scores could become an effortless second nature capacity (more on this in section 4.3). Using descriptive attitudes is only easier in the way using a rubric to mark an essay is easier. For experienced markers it might only get in the way, but for most people it is a simplification that makes the task easier.

Weightify, when processing the input information, can convert these attitudes into a numeric credence value. There are a few options for how Weightify then handles the conversion. Each attitude could simply be assigned a set value (e.g., a strong inclination that *P* = a credence of 0.8). However, we could treat them as floating values, because they do actually represent a range of possible credences (e.g., a strong inclination that *P* = a credence between 0.7 and 0.9). When Weightify processes the data, it does so many times (using the extremes for each credence and possible set intervals) and gives us a range of appropriate credences, rather than a specific one[73]. The credences are processed using numeric values, but Weightify's output could also be delivered using the descriptive attitudes. We could achieve something similar with expertise by selecting an expertise rank and letting Weightify convert that into a value. Adopting scales with clear categories makes the task of assessing expertise and credence easier both because we aren't trying to convert assessments into values, and because we don't have to make such detailed assessments. This also simplifies the data input process, which is now simply a matter of choosing the right box on a grid.

[73] It could still recommend a single credence as the best one – this range would have a normal distribution.

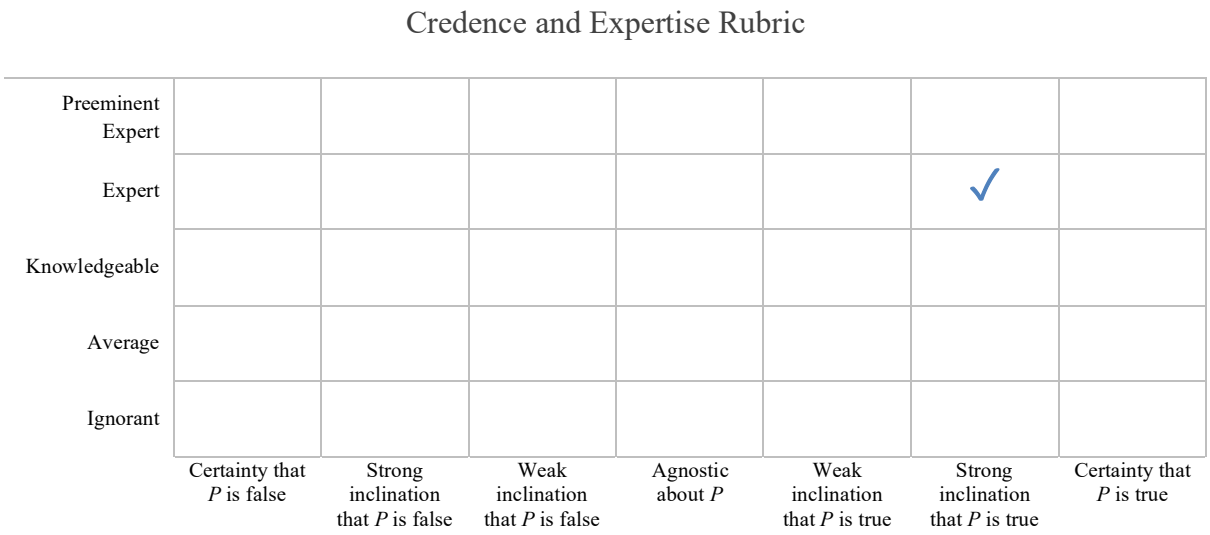

Fig. 12) Credence and Expertise One-Touch Simplified Input

This change may appear to cost the PEWV some accuracy, but I don't think that's actually the case. Firstly, the Weightify app can account for the fuzziness of these descriptive categories when it processes the data, to produce an appropriately fuzzy output. Secondly, if it is the case that we aren't that good at making fine-grained assessments, then we aren't losing accuracy by abandoning fine-grained assessments – those assessments were unlikely to be accurate to begin with. In any case, the level of granularity is flexible, and the user should aim to record as much detail as they are confident they can discern (all the way up to directly selecting fine-grained credence and expertise values).

At this point making an entry in Weightify happens very quickly compared to using the paper ledger. I need to open the app, select a belief to add to, tap on the grid to select a credence and expertise score, and name the entry. Aside from typing the name, that only took three taps! After the entry is made Weightify instantly processes the updated ledger and tells me what my new electoral assessment ought to be. With some shortcuts and a well-designed user interface the task could be made even easier, and with practice the whole process could be done very quickly.

But is the task easy enough? Beyond switching to coarse attitude assessments for expertise and credence values, I haven't discussed how we actually collect this information here. We could set down the path of externalizing data collection. For example, we could share our Weightify data online and build meta-entries that are collections of other people's ledgers,

but at that point the PEWV stops being a simple tool for resolving personal disagreement and becomes something larger, something outside the domain of this thesis. Thus, I will discuss ways to further reduce the cognitive burden of the collection module by simplification, rather than externalization, in section 4.3.

Ignoring the difficulties associated with the collection module for now, does Weightify reduce the cost of using the externalized PEWV enough? I think that it's certainly an attractive option for people who buy into the EWV. Almost all of the cognitive cost of the application is in data entry, and that has been made much easier than it was in the IEWV. Because Weightify processes the ledger for us, the cost of using the application also doesn't grow as we meet more interlocutors. While the application cost is static and fairly low (or front-ended if you want to set up a new belief and have to make many entries at once), Weightify can't escape the innate cost of being an external app we have to physically interact with.

For people who are already comfortable using computer programs as part of an extended mind this isn't much of a hurdle, but for others adjusting too frequently using an app like this could be too annoying or exhausting.

Even for those who are inclined towards using it, Weightify could not be used to resolve *all* cases of disagreement. Humans disagree about thousands of things, creating a record for all of these beliefs would be a massive undertaking. As it was with the IEWV, it is much more likely that we would only use the Weightify app in cases where we really care about our electoral assessment. These are cases where either the subject of the belief is very important to us, or where what we believe will shape an important decision. For example, I might be willing to use Weightify to arrive at an electoral assessment regarding big philosophical, religious, political, or economic questions, but ignore it for more trivial day-to-day disagreements.

4.3 The Simplification Approach

Despite the increased efficiency, the externalization PEWV (i.e., Weightify) has over the IEWV, it is unable to avoid the interfacing costs that come with using any external tool. Externalization yielded an application that is more accessible than the IEWV but is

nonetheless too cumbersome for use with unimportant (and more common) disagreements. The simplification approach, which I describe here, aims to rebuild the IEWV so that it doesn't require any externalization – neither for memory nor calculation.

The IEWV used parallel credence processing to avoid the sequence and stubborn peer problems outlined in section 2.6. Unfortunately, parallel processing *requires a ledger* to function properly when there are many beliefs to track, as was shown in section 3.2. The external ledger is too costly to maintain for day-to-day disagreements (even with Weightify's increased efficiency) so parallel processing will also need to be replaced. When discussing the problems with parallel processing (in section 3.3.1) I mentioned an alternative originally brought up by Gardiner, which I referred to as "weighted series processing". Weighted series processing involves processing peer credences in *series* (rather than all at once in parallel) but weighting the newest belief in proportion to the total of the beliefs that have factored into the electoral assessment so far. That is, if my electoral assessment is made up of three peer credences and I meet a new peer, I will conciliate the two positions (i.e., my electoral assessment and their personal assessment) with a weight ratio of 3:1. Now that we are working with expertise values this process can also be used in disagreements with epistemic superiors and inferiors. Only a small change is needed: the weight of my electoral assessment is now the sum of all of the interlocutors' *expertise values* that are part of the assessment. The weight of a new interlocutor's assessment is just their *expertise value*. That means that I only have to track an electoral assessment (the output of the EWV) and the total expertise that went into that assessment.[74] Thus, factoring more expert credences into my electoral assessment makes it more resistant to change, and the higher an interlocutor's expertise is the more their belief will affect the electoral assessment after processing. So far so good. If the total expertise in my electoral assessment was at a value of 3, and I thought my interlocutor's expertise was 5, then the epistemic weight ratio should be 3:5, even though my electoral assessment included a few beliefs. After applying weighted series processing, my new electoral assessment would have an epistemic weight of 8.

[74] Originally this would have been the number of peers that I had met, but as the view ought to be useful beyond peer disagreement is has been changed to expertise.

As with the other processing methods I have discussed, to avoid counting interlocutors more than once I should only process my interlocutor's personal assessment.[75] Thus, to make collecting personal assessments a bit easier for my fellow PEWV users, and to avoid the election magnetism problem, I should also track my own personal assessment and be able to present it to my interlocutors when appropriate.

Conciliating beliefs with the weighted series application and expertise allows us to generate collective assessments that don't violate the principles of the EWV without needing a ledger. All that I need to remember is what my collective assessment is, how much epistemic weight belongs to it, and what my personal assessment is. Remembering exactly how much expertise has gone into forming a collective assessment is an additional memory requirement over something like the SV, but it might not be too difficult to recall for anyone who is used to this process. For many of my beliefs I already have a sense of how much other people I've spoken to agree with me – this application just requires a more detailed and accurate version of that sense. I think a lot of the difficulty here comes from having to remember the epistemic weight as the numeric sum of the expertise values. In section 4.2 I argued for reducing the cognitive cost of assessing and recording credences and expertise values by changing how we think of them. Instead of the IEWV's numeric values, we could use more tangible descriptive attitudes. I think that we should do the same thing in the simplification approach. Doing so could make remembering this information easier, and it could make processing beliefs less taxing. Of course, in the externalization approach Weightify enabled us to ignore fine-grained assessments by converting the descriptive attitudes we selected into numeric values for us. In the simplification approach we don't have a machine to convert or process the information on our behalf, and so we will need to somehow process these descriptive attitudes. I will discuss how we could do this shortly.

This version of the weighted series application has (relative to the IEWV) small memory requirements. But, if we don't switch to using descriptive attitudes for expertise and credence assessments it also requires a lot of calculation: I must multiply my electoral assessment and my interlocutor's personal assessment by their respective expertise values, add the results,

[75] By failing to account for my interlocutor's interlocutors, and so on. This problem was originally discussed in section 2.6.4.

and divide by the new total expertise value to arrive at my new electoral assessment.[76] This calculation is going to be very taxing for most people – but the application can't require using a calculator! So, while this calculation is easier than the one found in the IEWV (at least after a few personal assessments are in the ledger), it needs to be made even less cognitively expensive.

The math used when I combine my electoral assessment with my interlocutor's personal assessment using the weighted series process allows me to combine them in the correct proportions. In a sense I am like a baker mixing ingredients. A novice baker, when mixing flour and water to make bread, might need to use a calculator to work out how adding another 100ml of water to his dough will affect the final loaf. But an experienced baker doesn't need to use a calculator or refer to a cookbook; without *really having to think about it* she knows how the additional 100ml will affect the loaf. The experienced baker has, over the course of her career, developed the ability to process these flour and water ratios using some heuristic, so she doesn't have to manually (or consciously) compute them. Following Gigerenzer, a heuristic is a fast and cognitively cheap method for making an inference – usually one that doesn't require computation or reasoning.[77] The baker's heuristic might not be perfect, but it works well enough and with such a tiny cognitive cost that she can employ it all day without fatiguing. With some work I think we could develop similar heuristics to avoid manually calculating the weighted series process.

Adam Morton argues that such heuristics can be a big part of what he calls an intellectual virtue.[78] Morton's intellectual virtues are very different from those discussed by Kelly in section 1.4.3.[79] Traditional conceptions of intellectual virtues, like those Linda Zagzebski talks about, tend to conceive of them as character traits with a positive normative value.[80] For example, it would be a good thing to properly develop the intellectual virtue of open-

---

[76] $\frac{(electoral\ assessment * its\ expertise)+(personal\ assessment * its\ expertise)}{sum\ of\ expertise}$

[77] (Gigerenzer, 2001)

[78] (Morton, 2012)

[79] (Kelly, 2005)

[80] (Zagzebski, 1996)

mindedness and thus become an open-minded person. These epistemic virtues tend to represent characteristics, character traits, or personality traits, rather than capacities.

This is not so for Morton's virtues as his outlook is more instrumental. If I don't care about a certain domain then I'm no better off with or without epistemic capacities that would only be useful in that domain. For Morton intellectual virtues are capacities that enable us to achieve certain intellectual goals.[81] That is, they are directed towards producing a certain outcome. Such virtues will often employ heuristics because they are cheap and effective tools for achieving these goals.

Whether we ought to give these capacities the title of epistemic virtues isn't a topic I'm going to delve into here. I am merely looking for a cognitively cost-effective method to reproduce the results of the IEWV. I think we can do that by developing heuristics that replace the cognitively taxing modules of the IEWV. Still, Morton's discussing of epistemic virtues is useful in this thesis because he tells us a bit about how we could come to acquire the capacities I am looking for.[82]

The best candidate heuristics for replacing the weighted series process are as unconscious or intuitive as possible because such heuristics are very fast and very cheap. They make use of what Daniel Kahneman calls System 1 processing (the type of fast, unconscious, semi-automatic thinking we use to, for example, solve 2+2=?), rather than System 2 processing (the slower, more deliberate and conscious thinking we use when, for example, solving a difficult math problem).[83] It is tricky to explain just how such heuristics would work, because they could be a little different for everyone, and are unconscious. In fact, because these heuristics are so user-specific and opaque, I think it is better to outline how they might be acquired than to attempt to explain what they might look like.

So how do we acquire such a heuristic? Say you have an electoral assessment that *P* is highly likely to be true based on disagreements with many expert interlocutors (giving this assessment a substantial expertise total). You then learn that another epistemic peer disagrees

[81] (Morton, 2012). This is attitude is also found towards heuristics in (Gigerenzer, 2001).

[82] (Morton, 2012, p. 69) see "A model for virtue acquisition" in Chapter 3.

[83] (Kahneman, 2011)

with you. Instead of using the IEWV or Weightify, you want to use a cheap heuristic to quickly gauge how this new disagreement impacts your electoral assessment. Ideally this heuristic would be unconscious, fairly accurate, and automatic (like the baker's). I think that the Weightify app discussed in section 4.2 could actually help us to acquire it. Guesses are a common form of heuristic. We might be able to develop a guessing heuristic to quickly process credences by guessing how new interlocutors affect electoral assessments and comparing our guesses to the results in Weightify. By repeating this process, we could train our intuitions to pick up on patterns in how interlocutors affect different electoral assessments, until we get to a place where the most common processes can happen pretty much automatically (similar to how I don't have to actually multiply 5 by 5 in my mind to know that 5x5=25, if a little bit less clearly defined). We become better over time by slowly building up a "library of stereotypes and associated reactions to them".[84] Perhaps Weightify could even have a dedicated training mode designed to help us quickly acquire these capacities.

We can imagine a heuristic method for quickly and cheaply performing every part of the EWV. Indeed, we might have multiple heuristics to reproduce the same modules in different circumstances.[85] Among other, more specific processes, there are potential heuristic replacements for:

- Assessing an interlocutor's personal assessment.
- Assessing an interlocutor's expertise.
- Computing the weighted series process.
- Recalling an electoral assessment.
- Recalling the expertise value of an electoral assessment.

Appealing to heuristics could also help explain what we are to do when collecting our interlocutors' credences and expertise values. If any of our interlocutors also practice the PEWV they will be able to tell us what their personal assessment is,[86] which gives us

[84] (Morton, 2012, p. 69)

[85] (Gigerenzer, 2001)

[86] As could a strong SV adherent.

something to measure how good we are at gauging these assessments without being explicitly told what they are. We could build up intuitive heuristics by talking to these people, attempting to gauge their credence, and then asking them what it is. Again, the hope is that over time this results in our System 1 processing becoming increasingly accurate.

Expertise is a bit more difficult because it is not clear how we could acquire a "correct expertise value" to test ourselves with. If I trust someone else to assess expertise better than myself, I could use their judgments as a yardstick to measure myself against. Such assistance might be available in some domains, but there is no final arbiter to tell is who truly is an expert. Thus, we must rely on our own judgments. I think using a heuristic method to assess expertise is actually very common (after all, we are lazy thinkers). We would just have to test these heuristics against our own more deliberate expertise assessment methodologies to assess their viability.

Acquiring the various capacities that would be of use in applying the PEWV doesn't necessarily require testing heuristics against an external measure. I think that having access to external tests is useful for making sure that we actually possess the capacities for a heuristic-based application of the PEWV, but not all heuristics are System 1 capacities. For heuristics that are just a simple rule we only need to be told about them. For example, while there were problems with using collective assessments in the IEWV, if we are already relying on heuristics the high epistemic value to cost ratio of a collective assessment outweighs the potential loss of accuracy that accompanies it. Rules like "try to find high expertise collective assessments about *P*" or "ignore low-quality personal assessments" make applying the view significantly easier by reducing the number of times we need to apply it, without sacrificing too much accuracy. I think it is a strength of the simplified PEWV that many of us already make use of heuristics like this every day.

As a general rule, heuristics won't perfectly replicate the expensive and careful forms of cognition they are replacing. They are only meant to be *good enough.* If we don't actually possess these capacities and try to use them, we are going to make a lot of mistakes, and our electoral assessment will not look anything like the assessment that would have been produced by the IEWV or by using Weightify. Even if use *fairly* accurate heuristic methods, repeating *slight* inaccuracies in recursive applications will result in our electoral assessments

straying from those recommended by the IEWV. However, there is a range of imperfection where the results would still be closer to an ideal application of the EWV than those produced by the alternative application (e.g., the standard EWV). Because the heuristics methods I recommend we acquire are so cheap they are well suited for resolving common and unimportant disagreements. It is only when we cross that line that the heuristic method stops being the most practical solution for day-to-day disagreements.

4.4 Conclusions

The two approaches to reducing the cognitive load of the IEWV to produce the PEWV produced very different applications. How close did either approach come to being a truly *practical* EWV?

The externalization approach led to the Weightify app. Depending on how we decide to input data into Weightify, it can produce exactly the same results as the IEWV while being much less demanding to use. Weightify automatically handles data processing, so there is no cost associated with that. Unfortunately, while Weightify does a lot to make recording credences and expertise scores easier, without the ability to read our minds there will always be a substantial cost involved in simply interfacing with the application (relative to thinking in our heads, anyway). For people who grow up externalizing parts of their mind this cost might not be as great, or they might be more willing to pay it, but for most people alive today I don't think something like the Weightify app would be used for everyday disagreements. When we really care about having the right electoral assessment to work with it would be a valuable tool, and that might be enough to count the application something of a success.

On the other hand, the simplification approach abandoned any reliance on external systems in favor of cognitively cheap heuristics. Switching over to a weighted series credence processing methodology that also factors in expertise reduced the memory requirements of the application enough that it shouldn't require an external record to be used properly. Appealing to heuristic methodologies for everything from working out what our interlocutors' credences and expertise values are to processing that information to arrive at a new electoral assessment result in an application that is cheap enough to be used frequently on day-to-day disagreements. It is a fair point that, as it is described here, this approach is really just a possible path to the PEWV, contingent on the user's ability to develop the

relevant heuristics– heuristics which have not been defined in any detail. Still, I think that by using the various training methods discussed in this chapter there is a good chance that we could acquire such heuristics. Until we do, we would know (through deliberate training) that we don't possess them, and thus that any attempt to use a heuristic-based PEWV application would be deeply flawed and untrustworthy. Actively training these heuristics is very important – if we don't, we have no idea how safe our application actually is. A peculiarity of this application is that most of the effort is front-loaded into the process of acquiring these virtues. I think that we are usually willing to pay front-loaded costs like this to acquire new abilities if we think they are worthwhile. However, as it stands exactly what this cost is unknown (i.e., how long and how much work is required to acquire these virtues?), and possibly unknowable, and this might prohibit adoption.

I think that while neither version of the PEWV presented in this thesis is itself a perfect practical solution to the PPD, taken together they provide a fairly comprehensive and accessible methodology. When compared to the standard EWV I think these features of both applications stand out as good reasons to think of them as improvements:

- Due to the inclusion of expertise, they can be applied to disagreements with epistemic superiors and inferiors, greatly increasing the practical utility of the theory.
- Due to how they process credences, they function without violating the principle of equal weight when exposed to disagreements with more than one person (see the sequence problem in section 2.6.3).
- Due to how they use personal and electoral assessments, they don't result in epistemic stagnation upon widespread adoption (see the election magnetism problem in section 2.6.2).

These features were also true of the IEWV. However, the PEWV applications manage to achieve all of this without being prohibitively costly to apply. Using Weightify carries with it a cognitive cost that is too steep to be used for every instance of disagreement, but it can be used far more readily than the IEWV could be. Weightify would be a useful tool for determining an electoral assessment for something we really care about, or for training heuristics for the simplified PEWV. The simplified heuristic-based application is, if properly developed, even cheaper than the standard EWV would be. After all, it wouldn't require any

manual computation. It is not as accurate as Weightify, but it should be good enough to be used for more common unimportant disagreements. I think that trading prohibitively expensive perfect results for pretty good cheap ones is a sensible move to make for limited agents like ourselves. PEWV applications are better than the standard EWV at reflecting the principles of the EWV in the electoral assessments they generate. Together they also form a framework for handling disagreement that we could actually use in our normal human lives. It is not an ideal framework – but it is a practical one that actually gets its job done.

***

Bibliography

Bentham, J. (1983). *Deontology: The Collected Works of Jeremy Bentham; Together with a Table of the Springs of Action and the Article on Utilitarianism* (G. Amnon (ed.)). Clarendon Press.

Carey, B., & Matheson, J. (2013). How skeptical is the equal weight view? In Diego Machuca (Ed.), *Disagreement and Skepticism* (Issue Section 4, pp. 131–149).

Christensen, D. (2007). Epistemology of Disagreement: The Good News. *The Philosophical Review*, *116*(2), 187–217.

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, *58*(1), 7–19.

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

Feldman, R. (2007). Reasonable religious disagreements. In L. Antony (Ed.), *Philosophers Without Gods: Meditations on Atheism and the Secular Life* (pp. 194–214). Oxford University Press.

Feldman, R., & Conee, E. (1985). Evidentialism. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, *48*(1), 15–34.

Frances, B. (2014). *Disagreement*. Polity Press.

Frances, B., & Matheson, J. (2019). Disagreement. In *The Stanford Encyclopedia of Philosophy* (Winter 201). Metaphysics Research Lab, Stanford University.

Gabbay, D. M., Hartmann, S., & Woods, J. (2011). Handbook of the History of Logic. *Inductive Logic*, *10*.

Gardiner, G. (2014). The commutativity of evidence: A problem for conciliatory views of peer disagreement. *Episteme*, *11*(1), 83–95.

Gigerenzer, G. (2001). Bounded Rationality - The Adaptive Toolbox: Chapter 3. In *Bounded Rationality - The Adaptive Toolbox* (pp. 36–50).

Gutting, G. (1982). *Religious Belief and Religious Skepticism*. University of Notre Dame Press.

Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.

Kelly, T. (2005). The epistemic significance of disagreement. In J. Hawthorne & T. Gendler (Eds.), *Oxford Studies in Epistemology, Volume 1* (Vol. 1). Oxford University Press.

Kelly, T. (2010). Peer disagreement and higher order evidence. In A. I. Goldman & D. Whitcomb (Eds.), *Social Epistemology: Essential Readings*. Oxford University Press.

Lackey, J. (2010). A Justificationist View of Disagreement's Epistemic Significance. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Social Epistemology*. Oxford Scholarship Online.

Matheson, J. (2015). The Epistemic Significance of Disagreement. In *The Epistemic Significance of Disagreement*.

Mill, J. S. (1863). *Utilitarianism*. Batoche Books.

Morton, A. (2012). *Bounded Thinking: Intellectual Virtues for Limited Agents*. Oxford University Press.

Skyrms, B. (2000). Choice and Chance.: An Introduction to Inductive Logic. In *Wadsworth* (4th ed.).

Sosa, E. (2013). Can There Be a Discipline of Philosophy? In *The Epistemology of Disagreement: New Essays*.

Weatherson, B. (2013). Disagreements, Philosophical and Otherwise. In J. Lackey & D. Christensen (Eds.), *The Epistemology of Disagreement: New Essays*. Oxford University Press.

Weisberg, M., & Muldoon, R. (2009). Epistemic Landscapes and the Division of Cognitive Labor. *Philosophy of Science*, *76*(2), 225–252.

Zagzebski, L. T. (1996). *Virtues of the Mind: An Inquiry Into the Nature of Virtue and the Ethical Foundations of Knowledge* (Issue 4). Cambridge University Press.