8 Paradigm Shifts and Group Belief Change

Haixin Dang

8.1 Introduction

What a scientific community holds to be its core beliefs changes over time. In the early twentieth century, the physics community rejected ether theories. Geologists slowly came to accept plate tectonics and continental drift in the 1960s. In the early 2000s, particle physicists revised the highly confirmed Standard Model in order to account for neutrinos having mass. How should we understand such belief changes within science?

A natural starting point for thinking about this question is Thomas Kuhn's seminal *The Structure of Scientific Revolutions* (SSR). In it, Kuhn lays out an account of how research communities come to accept a core set of theories and how these theories are supplanted by new theories over time. He describes these revolutions in science as cycles of normal and extraordinary science. It is hard to overstate how much Kuhn's ideas have influenced philosophy of science for the last sixty years. Practically every aspect of Kuhn's account has since been debated: from the concept of a research paradigm to theories of reference to whether Kuhnian revolutions occur as theorized. However, only recently has there been philosophical work done on how Kuhnian accounts of scientific change are a contribution to social epistemology. SSR not only highlights the social dimensions of scientific knowledge and the training of scientists as social epistemic agents embedded in complex social networks but also offers a convincing account of how scientific knowledge is *collective* knowledge. In particular, Wray (2011b) has offered a positive reading of Kuhn for his insights into the social epistemology of science. Wray points out that Kuhn later retracted the metaphor of the "gestalt switch" as a metaphor for paradigm shifts because a gestalt switch is an individual experience, whereas, for Kuhn, a paradigm shift is experienced by the scientific community as a whole (Wray 2011b, 15). Following this key observation, this chapter will explore how to understand Kuhnian paradigm shifts as a genuine social epistemic phenomenon, that is, as *group* belief change.

https://doi.org/10.1017/9781009122696.013 Published online by Cambridge University Press

Among social epistemologists, what a group belief is, or is not, and how it is formed have been well studied and debated (Gilbert 1987; Tuomela 1992; Mathiesen 2006; Bird 2010, 2014; List and Pettit 2011). However, much less is known about how group beliefs, once formed, may change over time (see Pettigrove 2016). There has been much philosophical debate over the nature of group belief in general (Hakli 2006). But there exist metaphysical worries over whether groups can have mental states and whether talk of a group belief is just metaphorical. Wray (2001) argues that group beliefs as studied by social ontologists like Gilbert are better characterized as group acceptance. For the purposes of this chapter, I will set aside the question of which cognitive attitudes can be properly attributed to groups. If one is particularly averse to theorizing about group belief, we can also make use of "group views." Belief changes are simply changes in views (see Harman 1986).[1]

Margaret Gilbert (2000) gave an account of group belief change by specifically picking out scientific change as the phenomenon of interest. She asks: "Is scientific change a matter of individual scientists changing their minds, or is it a matter of scientific communities changing their minds?" (p. 37) Gilbert has long defended an inflationary and non-summativist view of group belief. She believes that groups are agents in their own right, over and above the individuals (see Gilbert 1992). Group belief change cannot be just a simple matter of all or most of the individuals within that group changing their mind. Rather, group beliefs require something more, a joint commitment among the members of a group to believe or to change their beliefs *as a body*. Specifically, Gilbert argues that the joint commitment model of group belief can account for changes in scientific consensus. Weatherall and Gilbert (2016) have extended the joint commitment model of group belief to explicitly make sense of Kuhn's theory of scientific change. On Gilbert's view, a group belief is difficult to change as it would require a new joint commitment to be formed to replace the old one. A group's resistance to group belief change, as predicted by Gilbert, corresponds with Kuhn's observation that normal science is conservative and slow to process new evidence in ways that require a change of paradigm.

[1] It bears noting that "belief" may be a particularly reductive way to frame scientific attitudes. Scientists work with models and have complex methodological commitments that are difficult to capture as beliefs. While I will mostly be discussing how group beliefs change, these changes are also often entangled with changes in methodologies and modeling.

There have been several critical assessments of the application of Gilbert's joint commitment framework to science (Mathiesen 2006; Wray 2007; Rolin 2008; Fagan 2011). Wray (2001) and Mathiesen (2006) have argued that groups holding group beliefs as a joint commitment may not be epistemically rational, as group beliefs may not be updated appropriately in light of new evidence. For Gilbert, group beliefs are not strictly answerable to evidence, but rather what constrains group belief change is the joint commitment itself – it is a sense of obligation to the other members of the group to uphold what the group has already been committed to. This analysis threatens to render group belief change to be an arational or irrational process.

In this chapter, I will evaluate the account of scientific change given by Gilbert (2000) and Weatherall and Gilbert (2016). In particular, I argue that the *primary* normative constraint on group belief revision is the weight of the evidence being considered by the group, and not the normative constraints that arise from joint commitments. To understand group belief change, we need a better theory of how epistemic groups respond to evidence. I will only briefly sketch a positive view here, as working out a full-fledged account of group belief revision is beyond the scope of the present chapter. Instead, I will look back to Kuhn in SSR where he highlights the role of accumulating new evidence in the process of paradigm changes. While normal science tries to explain away new countervailing evidence and account for persistent anomalies, ultimately, a new paradigm is accepted when it can better account for the evidence and anomalies. Resistance to scientific change cannot be completely explained by the joint commitments among group members alone. Rather, Kuhn's account suggests that epistemic groups can be directly responsive to evidence.

This chapter will proceed as follows. I will first reconstruct the joint commitment model of group belief in more detail. Then I will critically assess the mechanisms proposed in Gilbert (2000) and Weatherall and Gilbert (2016) on how joint commitments constrain group belief change in science. Finally, I will sketch a positive view of how epistemic groups may respond to new countervailing evidence by looking to Kuhn's account of how anomalies are resolved in science.

8.2 Group Beliefs as Joint Commitments

Gilbert defends one of the classic inflationary accounts of group belief, which in other work she also terms "collective belief." In short, on Gilbert's non-summativist view, there is a collective belief that p if some persons are jointly committed to believing as a body that p. In contrast, a

https://doi.org/10.1017/9781009122696.013 Published online by Cambridge University Press

classic deflationary account of group belief is summative: a group believes that p when some or all of its members believe that p. Gilbert rejects simple summativism because she believes that there is an important sense in which a group can have beliefs over and above its members. That is, a group can have beliefs that are not held *personally* by the group members. Gilbert (1987) gives an example of a poetry reading group where a collective interpretation of a poem emerges from the discussions of the reading group's members. However, while individual members take the collective interpretation to be the view held by the group, it could very well be the case that no one individual member *personally* believes the collective interpretation. Divergence cases like this are taken by Gilbert to be evidence against various versions of summativism. Group beliefs are more than mere aggregations of member beliefs, more than just members voting.

According to Gilbert, "a group G believes that p" should be interpreted by the following conditions:

(i) A group G believes that p if and only if the members of G jointly accept that p.
(ii) Members of a group G *jointly accept* that p if and only if it is common knowledge in G that the individual members of G have openly expressed a conditional commitment jointly to accept that p together with the other members of G. (Gilbert 1987, 195, original emphasis)

Weatherall and Gilbert (2016) give a more concise definition of a group belief: "A group G believes that p if and only if the members of G are *jointly committed* to believing that p as a body" (p. 196, original emphasis). The authors also clarify that the idea of acting "as a body" is that "the object of the commitment, the thing being committed to, is to emulate as far as possible a single believer of the proposition in question, by virtue of the combined actions and utterances of the parties" (ibid.).

Central to both definitions is the technical concept of joint commitment. A joint commitment is, at the most basic level, a commitment of the will (Gilbert 1992). Gilbert (1990) specifically explains that a joint commitment is not a promise between parties to each do their part. Rather, a joint commitment is a basic notion that is at the very heart and foundation of human social relations. A joint commitment is created when each of the parties has expressed their personal willingness to be a party to it under conditions of common knowledge. Given a joint commitment, the parties to it are bound together and are subject to it until *all parties* concur in rescinding it. Importantly, while the parties are subject to a joint commitment, they have obligations and rights in relation to one another. When parties enter a joint commitment, they are entering a

particular normative situation where each is committed to act in a certain way in accordance with the content of the joint commitment.

Joint commitments ought to be distinguished from personal commitments. I may have a personal commitment to leave my office door open whenever I am inside to signal to students and colleagues that I am free to have discussions and answer questions. As a personal commitment, I can unilaterally decide to rescind my open-door policy at will. The philosophy department, as a whole, can decide to jointly commit to a department-wide open-door policy, perhaps as a result of a faculty meeting where the members openly express willingness to institute such a rule. This joint commitment would require that each member of the department leave their respective office door open whenever possible. On Gilbert's view, this newly formed joint commitment is not composed of personal commitments. The joint commitment is, rather, a normative situation where it is common knowledge between members that each of them has expressed their personal willingness to be committed in a specific way and that they all are so committed as a body. Under such conditions, if I fail to conform to a joint commitment, absent certain background understandings, then, I risk offending all of the other parties to the joint commitment. That is, if I personally rescinded from the collective open-door policy, I open myself up to rebuke from the other members of the group. Unlike in the case of a personal commitment, the joint commitment ties me to conform to the expectations of the group.

This is fundamental to understanding Gilbert's account. Reneging on a joint commitment is an offense because it is a violation of an *obligation* to the other parties, who have rights to one's conformity to the joint commitment. "This is a function of the joint commitment itself: a joint commitment obligates each party to the others to conform to it, at the same time endowing them with rights to such conformity" (Weatherall and Gilbert 2016, 197).[2]

In the case of collective belief, members of a group are jointly committed to believing that p as a body, under conditions of common knowledge. Once a collective belief has been formed, on Gilbert's account, it becomes extremely difficult for members who are party to it to disown it. It also becomes difficult for members to express *personal* skepticism toward the collective belief. Gilbert (2000) explains this in more detail:

If they speak contrary to the collective belief without a preamble, they will be regarded by both themselves and the others involved as acting out of line. More

[2] The normative force of a joint commitment can be quite strong. If one were to renege on a joint commitment, then the other members of the group can *rebuke* one. The members of the group can *demand* that one follow through with the joint commitment.

> specifically, they will have violated obligations they had to the others, and the others have the standing to take other measures responsive to this situation. These could take the form of anything from a mild rebuke to complete ostracism. Even if they do not speak without a preamble but make it clear that they are expressing what they personally believe, doing this only makes it clear that they are not personally in sympathy with the collective belief. That is liable to lay them open to suspicion. Hence *there is a significant initial cost* in every case of mooting an idea that runs contrary to a collective belief. (p. 42, original emphasis)

Note here that even in cases where a party of the collective belief expresses that they *personally* believe something contrary to the collective belief, this would open them up to suspicion by the group. There is a strong sense in which the collective belief *constrains and restricts* the personal speech of members.

A joint commitment to believe that p as a body does not strictly require a personal commitment from each member to believe that p, or for each member to personally act as though they believe that p. Rather, the members of the group are only required to act as separate mouthpieces of the group, upholding the belief in any context in which they are acting in their capacity as group members. This is why, for Gilbert, members must be careful to use a preamble to distinguish between personal speech and speech as a member of the group. Failing to use the preamble properly is a violation of one's obligations to the rest of the group. However, collective beliefs can be quite tricky, in that openly expressing that one does not personally believe the collective belief can still be costly. While it may not earn one a rebuke or exclusion, it still makes one vulnerable to suspicion from others.

Given this background on Gilbert's account, we can see now that once a group belief is formed, it is very difficult to change. It is even difficult for a member to openly speak against a group belief without some initial cost. A group belief change, in Gilbert's sense, will involve getting all parties to jointly commit to rejecting the previous belief and then to jointly commit to forming a new one. It is not enough that members personally reject the group view. Gilbert (2000) writes: "for a body of which I am a member to change its beliefs requires something akin to an agreement to stop believing that p together and to start believing that q instead, and this is not easy to achieve" (p. 46).

It would be useful here to contrast Gilbert's view with a deflationary account. On the summativist view, group belief change is much more straightforward. When enough members individually change their beliefs, the group belief change automatically follows. Group belief change would not require more than changing the minds of some sufficient number of individual members – a majority, a supermajority, or

perhaps everyone. On the other hand, for Gilbert, if a majority or even all of the individual members of a group believe something, it does not follow that there is a group belief, as it could be the case that they have not yet formed a joint commitment to believe it as a body. However, once the group belief is formed and the members are jointly committed in the right way, the group belief can remain stable no matter how individual beliefs may vary.

8.3 Group Belief Change and Scientific Change

Gilbert (2000) and Weatherall and Gilbert (2016) argue that scientific consensus can be properly interpreted as a group holding a collective belief through a joint commitment. Gilbert even argues that the difficulty of group belief change is a *virtue* of her view because she can successfully account for important sociological aspects of scientific communities. She notes two particular facts to be explained: (1) that collective beliefs act as "brakes" on scientific change and (2) that scientific change is often initiated by neophytes or outsiders.

Let us address (2) first. There is a prevalent public perception that major epistemic breakthroughs are made by young researchers. The Fields Medal, the most prestigious award in mathematics, is limited to recipients under the age of forty. Isaac Newton was twenty-three when he had his *annus mirabilis* and discovered – among many things – calculus and the law of universal gravitation. Albert Einstein was sixteen when he came up with one of his most important thought experiments: What would happen if he could chase a beam of light?

The idea that revolutionary change comes from either young scientists or those new to a field is also remarked on in Chapter VIII of SSR. Gilbert argues that young scientists are more likely to contribute revolutionary changes to science because they are not subject to the joint commitment in the same way as someone who is entrenched in the community – not being party to the scientific consensus allows one to go against it without fear of rebuke from members of the community. While Kuhn does not employ the technical notion of joint commitment, Kuhn (SSR-4) does echo a similar line of reasoning: "these are the men who, being little committed by prior practice to the traditional rules of normal science, are particularly likely to see that those rules no longer define a playable game and to conceive another set that can replace them" (SSR-4 90). However, Kuhn notes that while there exists a common perception that young scientists or outsiders play an outsized role in scientific change, whether or not this is true is ultimately an empirical question.

Wray (2003, 2004), through a survey of the historical record, showed that it is not the case that young scientists contribute more revolutionary research. Wray (2003) reports that a larger portion of significant scientific discoveries is, in fact, attributable to middle-aged researchers. Further, Line Andersen (2017) examines Gilbert's assertion that a scientist's outsider status can allow the scientist to more easily buck the consensus within an epistemic community. Andersen examines a case where outsiders to the mathematics community were able to make an important mathematical contribution. However, Andersen concludes that the facts of the case can neither confirm nor refute Gilbert's theory. We cannot make a judgment about whether the community's joint commitments played any role in preventing or bringing out the discovery. Given these empirical studies, it appears that Gilbert's claim – that scientific change is often initiated by neophytes or outsiders and this fact can be explained by the joint commitment model of group belief – has not been borne out.

Weatherall and Gilbert (2016), instead, focus on defending (1), the claim that collective beliefs act as brakes on scientific change. They argue that collective beliefs in Gilbert's sense can be analogous to a Kuhnian paradigm. Collective beliefs can regulate how the group functions, evaluates new evidence, and determines which projects are deemed appropriate for its members to pursue. Theoretical and methodological commitments associated with a Kuhnian paradigm may best be understood, they argue, as a set of foundational *collective* beliefs of the scientist within a paradigm. Becoming a full-fledged member of a scientific community requires one to jointly commit to what the community holds to be its core beliefs – to assent to the prevailing consensus and to work within the accepted paradigm. By entering a joint commitment, the scientist is then obliged to defend the consensus or otherwise face public rebuke or ostracism. Speaking out against the consensus – even with the preamble to signify that one is not speaking as a mouthpiece of the group – will still be personally costly. It casts doubt on one's commitment to the group and may signal unreliability to other members.

This analysis *appears* to echo much of Kuhn's analysis of normal science. Kuhn characterizes normal science as "puzzle-solving," that is, research based on established sets of scientific practice, with accepted exemplars and methods that can be fruitfully applied to new problems. When a scientist is working within normal science, their day-to-day research does not contradict the theories and methods that are foundational to that paradigm (see SSR-4 11). Even when faced with anomalies, or facts that seem not to fit the paradigm, scientists first try to explain away the anomaly with the tools available to them from within that paradigm.

Scientists do not immediately abandon a paradigm due to there being anomalies. Kuhn notes that "to be accepted as a paradigm, a theory must seem better than its competitors, but it need not, and in fact never does, explain all the facts with which it can be confronted" (18).

Normal science, then, tends to be conservative and thus slow to admit paradigm change. Countervailing evidence is often either explained away or ignored. Weatherall and Gilbert argue that the underlying explanation of this conservatism could be the fact that scientists within a paradigm are bound together by a joint commitment. They write:

> Even when evidence against a consensus is found by an individual scientist, one might expect it to be ignored, suppressed, or explained away by its discoverer, since such evidence would force a *psychologically unsustainable conflict* between the scientist's commitment to act in a certain way and his (or her) beliefs concerning the epistemic warrant for those actions. In the case of evidence contrary to a consensus that is made public, the scientific community, also, may ignore, suppress, or explain it away. For instance, it may be collectively affirmed that a crucial experiment cannot have been properly done, or it may be assumed that facts that are not at odds with the consensus can explain it, however implausible such an assumption really is. For these reasons, scientific communities *should be expected* to hold certain beliefs in the face of considerable conflicting evidence, to the point of being, or at least appearing, irrationally dogmatic or epistemically irresponsible. (Weatherall and Gilbert 2016, 201, my emphasis)

Note the mechanism the authors have posited as responsible for the group not being responsive to new evidence: members of the group are psychologically constrained by the normative force of the joint commitment itself. How does this mechanism work? Gilbert (2000) explains in more detail:

> It is important to see how participation in a collective belief can have consequences even for one's private thoughts, inhibiting one from pursuing spontaneous doubts about the group view, inclining one to ignore evidence that suggests the falsity of that view, and so on. Fulfilling the relevant commitment may not require that one personally believes what one's group believes. Nonetheless, it is awkward to say one thing and think another. This is so for a number of reasons. One may then resist following up certain ideas. One may – consciously or unconsciously – calculate that it is not worth doing so, since by making public the likely conclusions, even characterized as one's personal conclusions, one would risk a host of negative reactions to oneself and possibly even the breakdown of one's professional group. Knowing that this is something one is not prepared to do, one may avoid pursuing any relevant thoughts, even those one suspects of having something to them. (p. 45)

So, being party to a joint commitment will constrain what one can personally believe and personally do. Changing the collective belief would require all the members to jointly commit to a new belief.

However, individual members are, in a sense, psychologically constrained by the original joint commitment, so much so that one often discounts and resists the countervailing evidence, preventing members from coming forward to question the original joint commitment.

While Gilbert has always highlighted the importance of cases where members' personal beliefs may diverge from the group belief, there is tension in her account here. Even though an agent can form personal commitments, when their personal commitments clash with a joint commitment they are a party to, the agent is left in a psychologically "awkward" position. On Gilbert's account, the agent's obligation to uphold the joint commitment may be so strong that it overrides the agent's personal epistemic assessments. It is a consequence of Gilbert's view that even though a group member may personally recognize that new evidence ought to be accounted for, the group as a whole cannot respond to the new evidence. Surprisingly, once the agent is committed to being part of a body that believes that p, there appears to be some constraint on whether the agent can personally doubt that p.

Are epistemic groups this resistant to revising their beliefs, even in the face of convincing countervailing evidence? It bears mentioning that there is some intuitive appeal behind the idea that groups can be, or at least appear to be, irrationally dogmatic or epistemically irresponsible. For example, corporations often fail to change or adapt to evolving market conditions (Wright et al. 2004). There is now a growing literature on group polarization, a phenomenon where members of groups become more extreme in their beliefs after discussion with other like-minded members (Sunstein 2002; Del Vicario et al. 2016). While the dynamics behind these social phenomena are complex, these cases do show that groups often do fail, or at least appear to fail, to account for evidence in an epistemically responsible way.

For the remainder of this chapter, I will offer an alternative view of group belief change. In particular, I argue that the normative force of joint commitment in constraining group belief change is not primary. Ultimately, scientific groups are answerable to evidence. While scientific groups do seem resistant to certain kinds of new evidence, this resistance is not necessarily epistemically irresponsible. Rather, we need a theory of how a group can process and respond to new evidence, and we can look to SSR and Kuhn's analysis of revolutionary change in science as a starting point.

8.4 Kuhn and Group Belief Revision

Given Gilbert's arguments for the difficulty of group belief revision, it is surprising that groups can and do change their beliefs at all. While

normal science is conservative, periods of revolutionary science still occur. Gilbert (2000) acknowledges that the details of how a group changes its beliefs, especially a scientific group, are "worthy of study" (2000, 46). Weatherall and Gilbert (2016) apply the joint commitment model of group belief to explain how the string theory community is resistant to change and new evidence,[3] but the authors do not give an account of how an established scientific consensus eventually can be overturned. Because Gilbert claims that the normative force of the joint commitment will prevent group members from individually questioning the established consensus, this mechanism makes it mysterious how one may escape the joint commitment. I will argue here that the strength of the available evidence explains how both individual and group beliefs change.

Take a simple case: Mara is a member of a scientific group that has a collective belief that p, where all the members have openly expressed willingness to let p stand as the view of the group under conditions of common knowledge. Mara is committed to upholding p as a member of the group. However, she starts to encounter evidence against p during her research. Mara at first discounts the evidence, then she tries to explain it away in a manner that is consistent with p. She does this because of her commitment to holding p as a member of the group. Over time, Mara encounters more and more evidence against p. Mara begins to suspect that perhaps p ought to be rejected. Now, Mara is in what Gilbert had called a "psychologically unsustainable conflict" between her personal beliefs and the group belief. If she speaks out against p, she risks rebuke from her group members or, worse, complete ostracism. She might be prevented from getting more funding for her research or her papers may be rejected by journals. What should Mara do?

In Gilbert's telling, it seems that Mara has an obligation to keep believing that p. That is a consequence of being a party to a joint commitment. Mara may find it difficult to declare that she rejects that p even with a clear preamble to mark that she is speaking only for herself. If she does so, other members of the group will suspect her loyalty and exclude her from their future collaborations. So, Mara does not argue against her group and lets p continue to stand as the group belief. Gilbert

[3] The main claim in Weatherall and Gilbert (2016) is a conditional: *if* Gilbert's theory of group beliefs is correct, *then* the idiosyncratic features of the string theory community could be explained. Thanks to James Weatherall for highlighting this for me. In this chapter, I focus on assessing Gilbert's (2000) theory of group belief as applied to scientific change, but not to the string theory community in particular.

predicts that a widely held belief within a community will be difficult to change in this way.

But the evidence against p remains! Mara cannot ignore it indefinitely. Whether Mara ought to keep her obligation to her group has to be weighed against how strong the evidence is. Eventually, the evidence against p will be so strong and so persistent that Mara will be forced to speak out. She may have to argue with the other members of her group for them to also see the evidence against p. Mara's obligation to uphold and defend the joint commitment can go only so far. Her personal commitment to responding to the evidence in an epistemically appropriate way may be stronger than her joint commitment to hold p as the view of the group.

The most plausible way to understand the role of the joint commitment in constraining what agents can or cannot do is to view these special obligations that arise from a joint commitment to be only *pro tanto* obligations.[4] Mara lacks, all-things-considered, obligations to uphold the collective belief. She is also subject to other normative obligations, such as the norms of rationality or epistemic norms. Ultimately, if the evidence against p is strong and convincing, Mara ought to reject p, regardless of whether or not she has a joint commitment to uphold p. Once Mara has rejected the group view, she can then perhaps try to change the mind of the other group members. Epistemic norms must play some role in regulating how beliefs are formed and changed, at least, among members of scientific groups. A joint commitment can be broken, as it were, when the agents have other, stronger reasons to not follow through.[5]

[4] Elsewhere, in her work on joint action, Gilbert (1992) has claimed the agents have *pro tanto* obligations grounded in their joint commitments to perform their part of a joint action. However, agents may lack, all-things-considered, obligations to perform the joint actions, due to extenuating circumstances like side agreements or moral considerations (see also Kopec and Miller 2018). This should apply analogously to cases of collective belief as well.

[5] Analogously, for joint action, consider the following example. David and Nancy jointly commit to researching the efficacy of a drug together. This requires that David and Nancy coordinate their actions to bring about the study outcome as a group. Both Nancy and David at the beginning of their study believe that the drug is promising and deserves to be studied – perhaps they even jointly accepted that as a body. If Nancy fails to produce data according to their agreed-upon experimental protocol, David has the right to rebuke her for failing her obligations, and vice versa. Over the course of the collaboration, Nancy begins to suspect that the drug is harmful to patients and stops doing her part of the study. David still believes that the drug is promising and wants to continue the research as planned. Now, does David have the right to rebuke Nancy for not following through with their joint commitment to run the study? The normative force of the joint commitment does not automatically trump other competing normative considerations in this case. Agents are subject to many competing norms at the same time – both epistemic and moral. While a joint commitment gives rise to some specific obligations, agents can and do choose to act on other normative considerations. Nancy may fail the specific obligations she had toward the collaboration to follow through with a research plan.

If joint commitments are only *pro tanto* obligations that can be overridden by epistemic obligations, then members of scientific groups that have a group belief can still be responsive to evidence. Gilbert's view on scientific change deemphasizes how epistemic norms that govern how one ought to rationally respond to evidence can compel the members of a group to reject its original belief and adopt a new one. While normal science may be conservative in its evaluation of new evidence, scientists who are engaged in normal science are still responsive to evidence. When a paradigm has a demonstrated record of success, however, the strength of the countervailing evidence must be correspondingly strong – after all, extraordinary claims require extraordinary evidence. Extraordinary evidence is difficult to come by and alternative explanations must be evaluated before extraordinary evidence can be accepted as such. The fact that a scientific group resists revising their core theories in response to new evidence does not need to be primarily explained by the joint commitments of its members. Instead, the fact can be explained simply by the very nature of extraordinary evidence. What constrains the group from revising its core theories is not obligations to each other, but, rather, the group is constrained by evidence that is difficult to gather and difficult to interpret.

Another gap in Gilbert's account is how to explain why the members of a group jointly commit to upholding a group belief to begin with. How does a paradigm get established in the first place? Kuhn tells us that a paradigm becomes established because it is *successful*: it is good at generating new problems and good at giving a framework to study new problems. Scientists engaged in normal science are trained to solve outstanding puzzles using the tools and methods of the paradigm. A paradigm gains wide acceptance and becomes established because it is more successful than its competitors:

> Paradigms gain their status because they are more successful than their competitors in solving a few problems that the group of practitioners has come to recognize as acute. To be more successful is not, however, to be either completely successful with a single problem or notably successful with any large number. The success of a paradigm … is at the start largely a promise of success discoverable in selected and still incomplete examples. Normal science consists in the actualization of that promise, an actualization achieved by extending the knowledge of those facts that the paradigm displays as particularly revealing, by increasing the extent of the match between those facts and the paradigm's predictions, and by further articulation of the paradigm itself. (SSR-4 24)

Instead, Nancy can prioritize fulfilling a different obligation, the moral obligation to conduct ethical research that minimizes harm to patients.

Paradigms persist because they generate new problems for practitioners. A paradigm is successful not because it is complete, but rather because it invites practitioners to work out more of its details and to account for anomalies within its framework. Practitioners then adopt a paradigm and take up its commitments because of *what the paradigm can do*. That is, scientists have rational reasons to accept and work within a paradigm.

In contrast, for Gilbert, the reasons why agents choose to express openly a willingness to be jointly committed as a body are incidental. The joint commitment itself is the foundation of the social group, not the reason why the agents decided to form the joint commitment in the first place. Weatherall and Gilbert (2016) argue that theoretical and methodological commitments associated with a paradigm can best be understood as a set of foundational collective beliefs of a community of scientists. But why did *these* particular collective beliefs become paradigmatic and not another distinct set of collective beliefs? Gilbert is generally uninterested in giving an analysis of this question.[6]

Kuhn, however, does have a distinct analysis of why one paradigm triumphs over a competitor. The successful paradigm is better at solving puzzles and can account for more anomalies. When the paradigm comes to be accepted in a community, it seems that scientists sign on to the paradigm because they believe that the paradigm is good at solving problems that they have "come to recognize as acute." On Kuhn's description, adherents to a paradigm often have rational reasons to accept one paradigm over a competitor. Kuhn discusses this at length at the end of SSR Chapter XII, "The Resolution of Revolutions." In fact, in a later essay (ET 320–339), Kuhn explicitly lays out how competing theories can be rationally assessed. Theories should be accurate in their predictions within their domain, be consistent, be broad in scope, present phenomena in an orderly and coherent way, and be fruitful in suggesting new phenomena or relationships between phenomena.[7]

[6] For Gilbert, it does not matter what the reasons are because different members will have different and competing reasons to make themselves party to a joint commitment. Joint commitment is a general and basic fact about social groups. Members of one group can jointly commit to believing that p. Members of a competing group can jointly commit to believing that ¬ p. Gilbert is not interested in analyzing why some people believe p and while others do not. Instead, she is interested in analyzing how the joint commitment itself regulates the group and its members.

[7] Critics of Kuhn's account of paradigms often accuse Kuhn of introducing irrationality and relativism to scientific changes. But this is a misinterpretation of Kuhn's project (see Wray 2011b; Hacking 2012). While Kuhn's account of rationality in science may not satisfy those with more foundationalist leanings, Kuhn does not think scientists choose theories irrationally.

Once a paradigm becomes widely established and scientists are trained within it, scientists become very reluctant to reject it, even when faced with strong countervailing evidence. However, Kuhn's analysis of why practitioners of a paradigm take a long time to abandon an established paradigm diverges from the one given by Gilbert. No paradigm explains all anomalies. Much of the work done within normal science is trying to understand anomalies and account for them within the accepted methods and theories of the paradigm. However, members of the community openly discuss these anomalies as being unaccounted for within the paradigm. This is a signal that there is an open problem that would be a suitable research project within the bounds of the paradigm.

One of Kuhn's examples of how anomalies are handled within a paradigm is the discrepancy between Newton's original computation and the observed motion of the Moon's apogee, the point where the Moon is farthest from the Earth. This is a case of how an anomaly is resolved within normal science. For over sixty years after the publication of the *Principia*, physicists tried to resolve the discrepancy, without success. Newton's failed prediction regarding the motion of the Moon's apogee was well known and discussed. Kuhn even notes that during the time that the discrepancy remained unresolved, "there were occasional proposals for a modification of Newton's inverse-square law. But no one took these proposals very seriously, and in practice, this patience with a major anomaly proved *justified*" (SSR-4 81, my emphasis). Newton's inverse-square law proved too successful elsewhere to be abandoned, even when there existed important evidence against it. The community, in a sense, had faith that it could be explained. In 1749, Clairaut showed that the mathematics being used was wrong and that with simplified mathematics Newton's laws did accurately predict the observed motion. Kuhn uses this example to show that a persistent and recognized anomaly does not always induce crisis and cause the community to abandon a highly successful paradigm.

If the community had a collective belief in Gilbert's sense in Newton's laws, then we would expect those who speak out against it to be rebuked and ostracized. Perhaps we would expect that doubts cast on the Newtonian theory be formulated very cautiously and carefully, with a clear preamble and disclaimers. Yet, in this historical episode, it was the leading lights of the Enlightenment who openly discussed the problems facing the Newtonian paradigm.

For instance, d'Alembert had once remarked:

> The [Newtonian] system of gravitation can be regarded as true only after it has been demonstrated by precise calculations that it agrees exactly with the

phenomena of nature; otherwise the Newtonian hypothesis does not merit any preference over the [Cartesian] theory of vortices by which the movement of the planets can be very well explained, but in a manner which is so incomplete, so loose, that if the phenomena were completely different, they could very often be explained just as well in the same way, and sometimes even better. The [Newtonian] system of gravitation does not permit any illusion of this sort; a single article or observation which disproves the calculations will bring down the entire edifice and relegate the Newtonian theory to the class of so many others that imagination has created and analysis has destroyed. (quoted in Hankins 1985, 37)

In fact, Clairaut, d'Alembert, and Euler all publicly discussed that calculating the Moon's motion can only be approximated, which posed an important test for Newton's laws. In the summer of 1747, Euler explicitly acknowledges the need to correct the inverse-square law to account for the motion of the Moon in a paper read before the Berlin Academy of Science. Around the same time, d'Alembert also presents a paper to the Paris Academy and the Berlin Academy on the problem of the Moon's orbit. Clairaut, in late 1747, bolstered by Euler's and d'Alembert's earlier papers, openly suggests adding additional terms to the inverse-square law. This created much uproar among his colleagues. Note that in suggesting a modification to Newton's laws, Clairaut is openly questioning a core commitment of Newtonianism. Newton's laws had been so well confirmed on other matters that modifying the law just to account for lunar motion seemed unjustified and ad hoc. In particular, his peers criticized him because "Clairaut's additional term, specifically introduced to explain the motion of the Moon's apsides, implies that the laws that rule the motion of Earth around the Sun, the fall of an apple to the ground, and the orbit of the Moon are no longer one and the same" (Bodenmann 2010, 28). Euler rejected Clairaut's modifications on mathematical grounds (Bodenmann 2010, 29) – not because he took Clairaut to be a bad Newtonian. Surprisingly, in 1749, Clairaut finally found a way to reconcile the Moon's observed motions with Newton's laws, not through the addition of new terms but through a simplification of the mathematics used, thus confirming that Newton's laws were correct after all.[8] Clairaut turned himself from a heretic to the great defender of Newtonianism in the span of just two years.

Gilbert predicted that within normal science, anomalies would be ignored or explained away because members have jointly committed to upholding the core tenets of the paradigm. In this case, we see that

[8] For a full accounting of this historical episode, see Hankins (1985) and Bodenmann (2010).

anomalies were openly discussed and were actively worked on as open problems – a successful paradigm leaves some anomalies open for exploration. But as the Moon problem persisted, Clairaut, d'Alembert, and Euler all started to publicly question the inverse-square law in the 1740s. They did not suffer, as far as we can tell, a "psychologically unsustainable conflict" between their personal beliefs and the group belief. Rather Clairaut, d'Alembert, and Euler were eager to give a novel solution to the Moon problem to make their name. The result of their speaking out is not ostracism. Much of the rebuke Clairaut received challenged him on the merits of his proposal. It was, indeed, difficult for Clairaut to convince his community that Newton's laws needed to be modified in 1747. But the reason the community resisted can be explained by the previous success of the Newtonian paradigm. The resistance, in this case, to Clairaut's modifications was justified, Kuhn suggests, because there was a way to reconcile the calculated and observed motion of the Moon that confirms Newton's laws. The true mark of a successful paradigm is its ability to, eventually, account for the trickiest of anomalies. Clairaut's colleagues did not resist his modifications because they felt a sense of obligation to uphold what they have jointly committed to as a community.[9] Rather, they were skeptical that a paradigm with a long and proven track record ought to be abandoned without further investigation. The mathematical challenge of calculating the Moon's orbit – which essentially requires one to solve a three-body problem – is immense, so it was a rational response for the community to demand more proof and exploration of alternatives.

The reason scientists do not readily form new beliefs in response to new evidence is not that the previous beliefs arationally constrain members' ability to evaluate new evidence. Rather, epistemic groups are faced with several deep problems regarding the nature of extraordinary evidence, possible alternatives, and how to evaluate them. What it takes to rationally resolve these deep problems requires time and careful consideration by many group members. This accounts for why changing a scientific consensus is slow and difficult.

Gilbert's work on scientific change has been provocative and a fruitful starting point for thinking about the role of collective beliefs and evidence

[9] One may object to my characterization of the case here by arguing that Clairaut, d'Alembert, and Euler are all not British scientists, giving them a kind of outsider status. They were perhaps not the most committed to upholding the core paradigm. Defenders of Newton were much more dogmatic in Britain as opposed to the Continent. But even in Britain, the problem of the Moon's apsides was a well-known problem for Newton's theory. While not British, Clairaut, d'Alembert, and Euler all viewed themselves as Newtonians and working with the framework of Newtonian mechanics.

in theory change. While there are many striking similarities between Kuhn's observations about how normal science proceeds and Gilbert's theory about how collective beliefs work, there are several key features of normal science that the Gilbertian view cannot account for. In particular, Gilbert's view elides the important role evidence plays in group belief change. I have argued here that the *primary* normative constraint on group belief revision is the weight of the evidence being considered by the group, and not the normative constraints that arise from joint commitments. However, to fully understand group belief change, we will need a better theory of how epistemic groups respond to evidence. Kuhn's observations about the nature of normal science and anomalies are suggestive of a social epistemic account of how an epistemic group and its members can evaluate new evidence. Developing this account will be important for future work.