Making an Author in Radically Collaborative Research

Bryce Huebner, Georgetown University
Rebecca Kukla, Georgetown University
Eric Winsberg, University of South Florida

Collaborative authorship in science is the overwhelming norm, and articles with tens or even hundreds of listed authors are not uncommon. Across philosophy, there has been an increasing recognition that collaborative research and collaborative knowledge production are central epistemological practices and not just fringe phenomena. The philosophical literature on collective actions and intentions, collaborative knowledge, and group minds is exploding (Gilbert 1989, 1996; Huebner 2013; List & Pettit 2011; Rolin 2008; Thagard 2002; Tuomela 2002; Wray 2002, 2006). In spite of this focus on collective and collaborative phenomena, we think that questions about collaborative *authorship* raise distinct philosophical issues that have yet to received sufficient attention; sometimes these have been improperly assimilated to issues concerning collaborative knowledge and action. In this paper, we examine the unique difficulties inherent in establishing authorship in the context of collaborative research.

Collaborative research can be organized in a variety of different ways, and it can be performed in different material contexts. Not everyone involved in a research project is an author; indeed, we argue that people are sometimes listed as 'authors' on a publication when they should not be properly understood as authors at all. We examine three broad classes of research architectures using case studies. We consider collaborative research with multiple authors, collaborative research with a single author and many collaborators, and 'radically collaborative' research that is distributed widely over time and space, involves researchers from multiple disciplines, and typically results in publications with dozens of listed authors. (cf. Kukla 2012; Winsberg 2012; Winsberg, Huebner, & Kukla 2014). We hold that the first two types of collaborative research leave our basic understanding of authorship untouched. But we argue that an important range of research questions require radical collaborations that seriously challenge our understanding of and standards for authorship.

Authoring is not the same as knowing, though it cannot be understood independent of its relationship to the epistemic activities responsible for producing knowledge. Although *knowing* is not an inherently social or communicative notion, authorship is by nature social and communicative, even where it is not collaborative. Unlike mere knowers, (nonfiction) authors are *accountable* for the content and accuracy of the information they produce. An author must be able to justify and vouch for the truth of her claims, defend her products when challenged, and retract her claims when she cannot defend them in light of criticisms or new information. An author might make false claims, at least so long as she remains accountable and takes proper responsibility for them. Inquiries that are backed up by faulty processes of knowledge production and acquisition yield authors that are less reliable sources of information. But such failings are not failures in the status of *authorship*. Failing to know is not the same as failing to author.

There are currently no good strategies for understanding or ensuring accountable authorship in the context of radical collaborative research. We argue that neither the multiple-author nor the single-author model of accountability is viable in these cases. So, if the results that are reported by such collaborations are to be authored at all, they must be *group authored*. But we contend that a great deal of radically collaborative research has structural features that make group authorship impossible, thus making authorship *simpliciter* impossible in these cases. Even where there are good faith attempts to produce group authorship, major architectural barriers remain. This, we claim, poses concerns about epistemic security as well as our ability to interpret and decide when to trust the results that are reported in radically collaborative publications.

1. Catch and toss authorship and centralized control authorship

We begin by looking briefly at two kinds of collaborative research that do not pose a fundamental challenge to a traditional notion of authorship. First, we consider a type of collaborative research where a few authors work together to produce a single co-authored paper--much as we have done with this one. Borrowing a term from Andy Clark (1998, 105-6), we call this 'catch and toss' authorship. Second, we examine a kind of collective research that relies on a widely distributed form of information processing, but which retains a high degree of centralized control over the reported results. We use the first type of case to demonstrate a way in which research can be deeply social without compromising authorship; we use the second type of case to demonstrate a critical respect in which the distribution of information processing does not entail the distribution of authorship. With these cases in hand, we turn in the following section to a discussion of the kinds of radical collaboration that put more significant pressures on our ordinary understanding of authorship.

Let's begin with the kind of case most familiar to philosophers. When philosophers collaborate, they typically exchange ideas in a way that allows each author to retain epistemic authority over the resulting paper. In the most familiar form of philosophical collaboration, one author develops an idea (either in discussion, or as she writes her section of a paper). She then ‘tosses’ it to another author who ‘catches’ it, revises or extends it, and then tosses it back. This type of collaborative writing often consists in a long series of such ‘catchings’ and ‘tossings’, and the process is typically repeated until every member of the collaboration is happy to put her name on the resulting paper. These types of collaborations allow every author to be accountable for her own contributions. But, more importantly, since the writing and thinking are genuinely *interactive,* each author is likely to form a relatively clear sense of the structure of the overall argument. While there may remain some minor disagreements regarding what ought to be said, each author remains in position to vouch for the finished product. Each author is typically in a position to defend the claims that are produced through the collaboration when they are challenged, and each author is in a position to retract the claims that cannot be defended in light of criticism.

Of course, things are not generally this simple, and there is a different of ‘catch and toss’ process that arises when a project requires a unique technical skill or information that is only available to a specialist. In some cases complex statistical analyses are required that a lead author does not know how to carry out; in other cases, the relevant research may require a working knowledge of some range of scientific data that is only available to a person who has undergone special training. But, in these cases, a single author can do the work that is relevant to a particular section of a paper in a way that allows her to retain responsibility for her portion of the paper. For example, one author may ‘catch’ a data set collected by someone else, carry out a complicated statistical analyses, and ‘toss’ these results back to the lead author who inserts them into the paper.[1]

[1] Although ‘catch and toss’ collaborations fall reasonably comfortably within our traditional understandings of authorship, they can still create epistemic complications. When people come to a project with complementary forms of knowledge and expertise, which are drawn from different fields of inquiry, they must somehow find a way to negotiate the details of a collaboratively authored project. Precisely how this occurs differs from collaboration to collaboration. At one extreme, there are collaborations in which multiple revisions of a project are carried out, with an emerging paper being tossed back and forth until every collaborator understands and agrees to every one of the claims that are made. This yields a situation where every author fully understands the target phenomena, and can thereby be seen as an author in the traditional sense of the term. At the other extreme, there are collaborations in which every author simply accepts the claims that are made by collaborators, developing no further understanding of the target phenomena beyond their original field of expertise. This yields a situation where authorship is distributed across the members of a group, but where it is possible to recover structures of accountability by tracking the patterns of trust and acceptance within the group. As Hanne Andersen & Susann Wagenknecht (2013, 1892) note, collaborative scientific practice typically lies somewhere between these extremes. Even in small-scale collaborations "scientists in some areas learn from each other and in

Clear lines of responsibility are maintained in this type of co-authorship, and where the corresponding author is unable to vouch for a claims that is made in the paper she will know which collaborator has this capacity. This means that it is possible for collaborations to include specialists while still making sure that there is someone who is epistemically accountable.

In both of these cases, 'catch and toss' collaborations help to organize information in a way that functionally distributes the labor required for producing some purported piece of knowledge. The ideas produced through 'catch and toss' collaborations are often only possible because of the patterns of reciprocal feedback that arise between collaborators. However, each part of the co-authored paper is the product of an individual author who produces representations that can be 'tossed' back and forth. This yields a sort of joint authorship in which claims to knowledge are produced through a highly interactive process that leaves standard structures of accountability and responsibility in place.

But there are types of research, and accompanying forms of collaboration, that are too complex for this 'catch and toss' model to be practically possible. For instance, after the French Revolution, the *Académie des Sciences* famously instituted a metric system of weights and measures. They also stipulated that right angles would now be divided into 100 grades rather than 90 degrees.[2] Nautical navigation relied heavily on trigonometric tables, and using this new standard of measurement on the fly proved practically impossible. So, the *Académie* instituted the *Bureau du Cadastre* to construct new trigonometric tables using grades rather than degrees. The director of the *Bureau*, Gaspard de Prony, was enamored of Adam Smith's idea of 'distributing labor', and he hired a team of human computers to carry out the arduous task of constructing the trigonometric tables. The vast majority of these computers (around 90 of them) "were former servants and hairdressers, who had lost their jobs when the Revolution rendered the elegant styles of Louis XVI unfashionable or even treasonous" (Grier 2005, 36). As such, they had no mathematical training beyond the basic abilities required to add and subtract. So, a team of eight 'planners' supplied these computers with worksheets that allowed them to carry out simple calculations; the planners then took differences between the calculated values to spot-check for errors. Finally, a small number of trained scientists figured out the appropriate formulas to be passed down to the planners. All of the relevant

others remain epistemically dependent." So, the social epistemology of even small-scale 'catch and toss' collaborations can be more vexed than we have suggested here.

[2] Our discussion of De Prony in this paragraph follows Grier (2005, 34-38). For an interesting discussion of this research in another context, see Shieber (forthcoming).

calculations were eventually passed along to de Prony, who put together a nineteen-volume manuscript that included all of the tables.[3]

Structurally speaking, this kind of research shares much with the 'catch and toss' collaborations discussed above. However, in this case, the human computers became nothing more than sources of information for de Prony's mathematical investigations. The organization and distribution of intellectual labor instituted clear lines of responsibility, such that de Prony (and perhaps his team of trained scientists) retained accountability for the resulting tables. The human computers became mere tools, to be used in the service of a project carried out by a few highly trained and highly intelligent scientists. Because of the organization of this research, de Prony (and perhaps the trained scientists) remained epistemically accountable. These authors could vouch for the results of this massive collaboration; they were the ones who were epistemically accountable for producing accurate tables, defending them if challenged, and revising them if necessary. The relevant structures of epistemic responsibility and accountability were maintained in the *Bureau du Cadastre* to allow for the distribution of collective research while leaving traditional notions of authorship intact. In effect, the text was still single-authored, even in the face of distributed information processing, because one person retained centralized control over the research process, including its methodological standards and implementation. While many people participated in the production of knowledge, only one person had the status of the author of the document communicating that knowledge.[4]

2. Genuinely distributed epistemic labor

In the type of case we just examined, labor is distributed, but there is no need for lower-level agents to exercise special epistemic skills or judgment, because the author who retains centralized control can provide lower-level agents with detailed and specific instructions about how to handle every problem that they might encounter. We now turn to a set of cases involving distributed epistemic labor, in which the primary reason for involving a large number of actors is that no single actor *can* possesses the relevant knowledge and skills that are necessary to *produce and sanction* the desired knowledge claims. In such cases, multiple actors must exercise special epistemic skills and judgment. In the types

[3] Unfortunately, the manuscript was never published because the publisher went bankrupt, and the Napoleonic government had no interest in publishing the volume.

[4] Strikingly, a similar model of highly distributed, centrally controlled research was instituted at the Oak Ridge National Laboratory as a way of isolating isotopes of Uranium-235. For a detailed discussion of this case see Winsberg, Huebner, & Kukla (forthcoming).

of collaborations that interest us here, there is no one who understands the role that is played by every researcher, and there is no one who knows what everyone else has contributed to the project; there is no one who has even testimonial knowledge that the other researchers are trustworthy; there is no one who has command over how the various pieces of the study fit into a coherent whole.

We argue that the reports that result from these types of radical collaborations cannot be understood as multiply authored, as in the catch and toss model, or single-authored while involving distributed labor, as in the centralized control model. Instead, if they have any kind of author at all, it must be a group author. Unfortunately, as we will see, the architecture of radically collaborative research is not typically of the right sort to produce a group author. Nor is it at all clear, given the real pressures and constraints that such research faces, how it could be restructured so as to enable group authorship.

2.1 Modeling in Climate Science

No case exemplifies radically distributed epistemic agency better than climate science—especially when it involves massive, modular, and highly complex *coupled atmosphere-ocean global climate model* (AOGCM) simulations, such as the National Oceanic and Atmospheric Administration's GFDL CM2.x Model.[5] CM2.x contains over a million lines of code, with over a thousand different parameter options. Its novel component modules were written by over one hundred people. These modules are constantly changing and they involve hundreds of initialization files that contain incomplete documentation. Loading input data into a simulation run takes over two hours. It takes weeks to produce a model run out to the year 2100, using more than one hundred processors running in parallel; and it takes months for them to reproduce thousands of years of paleoclimate data. If you store the data from a state-of-the-art GCM every five minutes, it produces tens of terabytes of data per model year. But these data must be collected relatively quickly to be of any use. Policy makers want immediate answers to near-term possible climate changes, and we would like to know what characteristics the climate is likely to have in 2050 sometime before 2050. This is part of the reason why these climate models must reflect the work of hundreds of researchers working at different times and in different physical locations. Furthermore, the time-indexed nature of the research makes it inherently unreproducable.

[5] Details concerning the CM2.x model come from (Dunne 2006).

The sheer size and complexity of such models make it clear why their construction and sanctioning must involve a vast army of specialists. Their production and use requires expertise in climatology, meteorology, atmospheric dynamics, atmospheric physics, atmospheric chemistry, solar physics, historical climatology, geophysics, geochemistry, geology, soil science, oceanography, glaciology, paleoclimatology, ecology, biogeography, biochemistry, computer science, mathematical and numerical modeling, statistics, time series analysis, and more. Furthermore, no GCM is built from the ground up in one short surveyable unit of time. They all rely on assemblages of methods, modules, parameterization schemes, initial data packages, bits of code, coupling schemes, and so forth that have been built, tested, evaluated, and credentialed over years or even decades of work by climate scientists, mathematicians, and computer scientists of all stripes.[6] This yields a sort of "fuzzy modularity" in these models (Lenhard & Winsberg 2010). In current atmospheric GCMs, we find a deliberately modular architecture that consists of coupled models, which are themselves comprised of numerous interactive submodels (e.g., modules for general circulation of the atmosphere, cloud formation, dynamics of sea and land ice, and effects of vegetation). The interactions between these models determine the global dynamics in simulation runs, and interactions can be quite complex.[7] In a modular system, individual models (ideally!) act as interchangeable but highly networked modules. The results of such modules can first be gathered independently and then synthesized. However, in current GCMs, the data is continuously exchanged between modules during the runtime of the simulation.[8] Thus the dynamics of the model can only be understood as the complex result of the interaction between these modules, each of which is infused with the judgments made by numerous different scientists, on the basis of their own assumptions about what the right way to parameterize is. To say that such models have only a sort of fuzzy modularity is to note that the interactivity and

[6] There has been a move, in recent years, to eliminate "legacy code" from climate models. Even though this may have been achieved in some models (this claim is sometimes made about CM2), it is worth noting that there is a large difference between coding a model from scratch and building it from scratch, that is, devising and sanctioning from scratch all of the elements of a model

[7] To take one example, the coupling of atmospheric and oceanic circulation models is recognized as a milestone of climate modeling. Both components had an independent modeling history, including independent calibrations of their respective performance. Putting them together was a difficult task because the two submodels interfered dynamically with one another. For an account of the controversies surrounding early coupling, see Shackley et al. (1999); for a brief history of modeling advances, see Weart (2010).

[8] Because data are being continuously exchanged one can accurately describe the models as parallel rather than serial in the sense discussed in Winsberg (2006).

the balancing of approximations prevents the system from being broken down into separately manageable pieces.

The problems for authorship that are evoked by this type of fuzzy modularity are exacerbated by the fact that epistemically salient decisions in climate modeling must be sensitive to a continuous, dynamic flow of information through the model. The collaborators must carry out continuous deliberative adjustments in light of new circumstances and new information, and this must frequently occur without the time or means to consult with other members of the research team. This is why the operation of each module must rely on a mixture of principled science and decisions about parameterization. Climate modeling involves literally thousands of methodological choices. Many crucial processes are poorly understood; many compromises must be made in the name of computational exigency; and so forth.

But no methodological choice can be defended in a value vacuum. When one asks, “Why parameterize this process rather than try to resolve it on the grid?” or “Why use this method for modeling cloud formation?” the answer is rarely “because that choice is simply objectively better than the alternatives.” Most of these choices are better in some respects and worse in others, and the preference for the one choice over its alternatives always reflects the judgment that this or that respect is more important. Some choices will result in more accurate predictions with respect to one variable (e.g. tropical precipitation), while its alternative will be more accurate with respect to another (e.g. polar ice). Some choices will increase the probability of finding a certain degree of climate variation; other will do the opposite. So any rational reconstruction of the history of climate science (if it existed) would have to mention each of these predictive preferences at pain of making most choices seem completely arbitrary (Biddle & Winsberg 2009, Lenhard & Winsberg 2010, Winsberg 2010).

Meanwhile, climate experts, in light of the individually limited role that they play in the socially extended activity of building climate knowledge, can only arrive opinions about the future of the climate in ways that are fundamentally mediated by the complex models that they build. And they are incapable of sorting out the ways in which past methodological choices made by other scientists--scientists whose expertise they don’t entirely share--are influencing, through their entrenchment in the very models that mediate their inferences, the conclusions that they deliver to policymakers.

No single person is a position to offer a rational reconstruction of a climate model. Too many of the methodological choices are buried in the historical past under the complexity, distribution, and historically path-dependent character of climate models. The various local standards and values employed in climate science lie in the nooks and crannies (Winsberg 2012). They might very well have been opaque to the actors who put them there, and they are certainly opaque to those who stand at the end of the long, distributed, and path-dependent process of model construction. There is no one who could, even in principle, be held accountable for the claims to knowledge that are produced using CM2.x. To put the point another way, there is no person who has the requisite expertise to play the role of de Prony in these modeling activities, and accordingly there is no single person who can legitimately be treated as the author of these claims.[9]

2.2 Multisite Clinical Trials in Biomedical Research

Multisite clinical trials, the purported gold standard of biomedical research, create an analogous pattern of difficulties. In many cases, biomedical research takes the form of geographically distributed projects that involve many teams of researchers, each of which uses nurses, lab techs, students, departmental administrators and many others to help implement the study and keep track of data. Papers with dozens of authors, often from several continents, are the norm. In many cases, no one has the expertise and training to design and implement every part of the study. Such research often draws on a wide variety of disciplinary expertise, including biostatistics, cell biology, oncology, and immunology, among many others.[10] These research projects are so dispersed

[9] Contrast this with even a very complex example of 'catch and toss' authorship: Timothy Gowers and Michael Nielsen (2009) report the results of a recent case collaborative research in mathematics known as the Polymath Project. The aim of the project was to find an elementary proof of the Hales-Jewett theorem. Over the course of 37 days, 27 mathematicians contributed approximately 800 substantive comments to a collaborative discussion, yielding 170,000 words of text, and an elementary proof of the Hales-Jewett theorem. Contributions to the proof were made by world-renowned mathematicians and university professors, as well as high school math teachers. Furthermore, contributors were active in the discussion to very different degrees, and at different points over the course of the collaboration. This case provides a nice set of insights into the operation of 'catch and toss' collaborations. The discussion reveals that ideas are often proposed, then adjusted, and sometimes discarded over the course of the collaboration. But perhaps more importantly it reveals that advances in understanding often result through the "aggregation and refinement of many smaller insights" (Gowers & Nielsen 2009, 880). This process generates practical questions about authorship. For example, how should credit be assigned when a contributor offers a single insightful contribution, and how should it be assigned when a contributor is prolific, but not particularly insightful? But a working record of the collaboration remains available for outside scrutiny. If there were a mistake in the proof, it could be tracked down, and it could be determined how mistakes were introduced, and how they were perpetuated in the project. This may not be an easy task, but at least it would be possible.

[10] The December 2011 issue of the *New England Journal of Medicine* includes five "original research" articles, with 27, 6, 29, 19, and 13 authors respectively. All but one of them is written on behalf of a much larger research group, and links to a supplementary appendix listing hundreds more

and so multidisciplinary that no one can have access to more than a small corner of these data. Furthermore, no one can have expertise in all of the on-the ground skills required to administer tests, read results, recruit subjects, etc. It is unlikely that the biostatisticians and immunologists fully understand each other's contributions, or that researchers in France know exactly how subjects were recruited in Taiwan, and so on. This is decentralized, distributed collaborative research on a massive scale – a scale frankly unthinkable until recent developments in communication technology.

The problem is not merely that there is no person or small group of people that could have the time or the expertise to run the study on their own. The data in biomedical research, as in climate science, are often time and place dependent, and questions about clinical effects are often context-specific. Policymakers frequently need to know immediately whether to close schools in response to potential epidemics, and there are sometimes political pressures to approve a drug or remove it from the market as soon as possible. Because medical interventions often work differently in different populations living under different social and material conditions, it common for a research question in biomedicine to *require* trials in far-flung locations, which must be conducted under messy and chaotic conditions, in the face of communication challenges. Additionally, since bodies change quickly, and since harms and benefits from a treatment often evolve over time, the time frame of a study (how quickly it must be completed, how close together in time the collection of different data sets must be) is often fixed by the research question. In other words the multi-site distribution of studies along with their rapid time frame is often essential to the research design. This, again, makes these studies inherently unreproducible.

Unlike centralized control cases, researchers involved in this sort of radically collaborative, distributed research must often make methodological adjustments on the fly in response to noncompliant research participants, unforeseen barriers to implementation and communication, surprising side effects, and so forth. It is hard to know in advance what methodological uncertainties, judgment calls and choice points collaborators will face. These adjustments may be made differently at different sites, and in different ways at different stages of the research process. Whereas De Prony's computers repeated the same task under controlled conditions, these researchers must draw on their own expertise to make on-the-spot decisions about how to proceed under imperfect and sometimes surprising conditions. For example, when researchers encounter participants who are

---

collaborators and participating investigators. All five articles list authors from various countries and various institutions within each country, and each contains authors with diverse disciplinary backgrounds. Other issues display roughly similar authorship patterns.

partially compliant, who disappear half way through a study, who display ambiguous symptoms, who are difficult to communicate with, and so on, they must still decide whether to use data from these participants or not. But this requires judgment calls about how much to bend protocol to get through the study.

For all these reasons, there is often no one involved in the research who has more than a loose idea who is doing what parts of the sprawling project, using what methods, or how accurately. Whatever being an 'author' comes to in the context of climate science or multisite clinical trials, it shares little with the traditional understanding of scientific authorship. Indeed, actual practices of assigning authorship make no pretense that authorship tracks responsibility for the production or the justification of the results. There is no presumption that the actual writing of the article, which many scientists see as extra-scientific busywork, will be performed or even overseen by the lead 'authors'. Authorship constitutes an institutional status; it does not represent a specific form of epistemic labor. Authorship is assigned and ordered according to whose grant money was used, who runs the lab, who has tenure, who needs a job, and so forth.

Often this reality is stubbornly denied by the researchers themselves. For instance, the many authors of "Boceprevir for untreated chronic HCV – A Randomized Clinical Trial" try to foreclose these worries by insisting that they have attained a more traditional authorial status. Their article states, "Each author vouches for the fidelity of the trial conduct to the protocol and the completeness and accuracy of the results and data analyses" (Poordad et al 2011). Unfortunately this is radically implausible, if taken as an epistemic claim. The article has 15 authors in four countries employed by 11 academic institutions plus Merck. It is presented on behalf of the SPRINT-2 investigators. The 68 page appendix lists 173 principle investigators from 11 countries who participated in the study, each of whom presumably relied in turn upon a wide variety of support staff to actually implement the study, keep track of the data, and communicate with other participants and with Merck. With the number of people playing some role in the path from study design to implementation to publication running in the hundreds or even thousands, the idea that *each* of the 15 authors could know that *everyone* involved was competent and reliable, that they adhered to the protocol in the same way, that the data set was complete and accurate, etc. is absurd. In radical, distributed biomedical collaborations, testimonial knowledge securing each part of the study is impractical.

3. Neither catch and toss nor centralized control

A *prima facie* response to our argument so far is to admit that radically collaborative science articles, as currently produced, have no author, but also to insist that scientific practice could be reorganized so as to bring it under the catch-and-toss or the centralized-control models. In that case, climate science, biomedical research, and similar massively collaborative, distributed, interdisciplinary sciences would differ from our original examples only in size and not in deep epistemic structure. We contend, however, that material constraints on this kind of research and the complex normative pressures it is under make it hard to see how to reorganize it in these ways. There are multiple barriers to the effective production of authorship in climate science and biomedical science that are not so easily overcome.

Consider first the attempt to assimilate authorship in radically collaborative science to catch-and-toss authorship. Even if there is no one person who can recreate a study, or even understand all its parts, it is not obvious why this is different from relying on a friend to run your statistics. One response to these worries might be to suggest that we reform the organization of radically collaborative science so as to allow transparency and coordination of methodological choices and standards. We need to ensure that information about the pressures that shape various parts of the study is readily available, on this line. Furthermore, we need mechanisms that guarantee that each author can have confidence in the competence and honesty of each other author. If we do this, we could, the argument goes, turn radically collaborative research into a huge catch-and-toss project. All of the authors could take responsibility for the whole study, because they have testimony-based confidence in all its parts and could hold their co-authors accountable for their interest-driven choices if needed. They understand what each of their collaborators contributed, know that each of their collaborators is competent and trustworthy with respect to her part of the whole, and can explain how those parts fit together. Legitimate collaborations would then have an embarrassment of riches: multiple authors, not just one author or no author at all.

In the catch-and-toss model, research can have multiple authors because they understand each part of the project and understand how they fit together, or because they understand why each researcher is trustworthy and accountable when it comes to her own part of the whole. But this requires the ability to know how all of the pieces of the research fit together into a coherent whole. This is not possible in the radical collaborations like those we discussed above. No researcher can vouch for the totality of the decentralized, spatially and temporally distributed epistemic labor. Radical collaborations draw on different kinds of expertise, and respond to different kinds of pressures that arise as a result of local context. There is no way for a climate scientist working with CM2.x to know the role of every action, value, and decision that went into building the model; there is no way for any individual to personally vouch for the fact that every contributor was competent and reliable in making his choices. Perhaps more importantly, even assuming the competence of everyone involved, there is no one who can know that those choices cohere with the methodological choices made by others who played a role in producing this model. Some of the history of CM2.x reaches back decades, and involved people and situations that our hypothetical single climate scientist has never even heard of.

The reasons why radically collaborative research can't be brought under the centralized control model are perhaps more interesting. Whereas de Prony relied on widely distributed epistemic labor, he was in charge of designing and implementing the project. He established the methodology of the study, relying on information processing that was largely mechanical and that yielded clear data that just needed to be assembled at the end of the production line. De Prony's computers did not need specialized epistemic skills, nor did they need to make their own methodological choices or establish their own epistemic standards; they simply executed the relevant functions. De Prony did not need to trust the expertise of his computers because they were not called upon to exercise any.

In other words, De Prony faced an engineering problem rather than an epistemic problem. As long as there were structural mechanisms in place to guarantee that his computers executed their function sufficiently reliably, and as long as he could count on relatively reliable information flow from the margins back to the center, his reliance on others to carry out the epistemic labor was not interestingly different from relying on any kind of data-collecting or data-storing instrument. De Prony was the only one who was establishing methodological practices and epistemic standards, and he was accountable for the justifiability of the methods and the accuracy of his results (including his method for ensuring the reliability of his 'tools'). This is the primary reason why he was in a position to take authorial responsibility for the representation of the relevant results.

Crucially, radical collaborations yield an epistemic problem rather than an engineering problem. In the cases that we have discussed, many different people are engaged in substantive epistemic labor; they are called upon to *use their judgment* and to *establish their methodology*, often in response to real-time messy pressures. This is the main reason why there cannot be someone who ensures the reliability of every decision and who collects and coordinates the information. No one has the relevant expertise to ensure that the epistemic practices and methodological standards used by everyone in the study are reliable. And lack of confidence in *reliability* isn't the only problem; since they are required to exercise judgment and expertise, the members of a radical collaboration cannot be treated as mere information processors. Their methodological choices and judgments contribute to the overall justificatory story in a way that is lost in simple attempts at collating results. Hence, no matter how much we work on centralizing and coordinating collaborative research projects, and no matter how much we improve transparency and information flow, there is no one who is in a position to play that authorial role of someone like De Prony.

This last point raises a more general and deeper problem: we cannot assume that different disciplines each have fixed methodologies and set epistemic standards. There are reasonable professional disagreements over how to proceed, but in any given situation, methods and standards of justification depend on the goals of the inquirers and the local pressures they face. If there were universally established, fixed methods standards for each type of research, we could set up something similar to the central control model; we would just need to create structural mechanisms that would guarantee that researchers would properly adhere to these methods and standards, along with reliable methods of information flow and coordination. In this case, the researchers could, in effect, function as *discipline-specific machines* whose results could be interfaced in a modular fashion. But it is precisely because this is not the case that we need researchers who have specialized epistemic skills and judgment; this is the reason why no central coordinator can establish in advance what everyone's proper standards should be. We mentioned that these sorts of problem arise in climate science above, but to make the generality of this issue clear, it will help to consider two additional examples.

First, it is standard for the principal investigators on a clinical trial to employ a dedicated statistician. Often, in an industry-funded trial, the statistician will be an industry employee. The lead 'authors' may or may not be in a position to check over the work of the statistician, looking for glaring errors. But the whole reason why research projects bring on a statistician, rather than using some sort of statistics software, is that there are substantive, nontrivial choices to be made among different kinds of statistical analyses. Running the statistics on a study is not merely a mechanical procedure, but rather requires subtle sensitivities to the project and pressures at hand. The PIs don't have the skills to make these subtle judgment calls, or they would just do the analysis themselves. But when research is radically collaborative, there is no clear way to guarantee that the goals that shape even a totally competent statistician's choice of methods and standards coordinate with the goals of other researchers. This is especially clear (but not only so) in the case of industry-employed statisticians working on privately funded studies, who have a clear motive to choose a method of analysis that is favorable from a marketing perspective. 'Authors' often have no mechanisms for ensuring that the goals of the statistician, and her understanding of the questions and pressures at stake in the research, coordinate with those of other researchers.

A different kind of example comes from local pressures that are opaque to spatially distributed collaborators. Consider an investigator who is on the ground collecting data on vaginal macrobicide and HIV transmission in a remote African village. It is likely that she will need to make many judgment calls as to whether a given subject was compliant enough with the protocol or giving reliable enough reports to be included in the study. But this is tricky business where there are cultural barriers, populations with low scientific literacy, and cultural forces and power relations that impede free and informed talk about sexual behavior. There is no single right answer that can be mechanically generated as to when a subject's data should be included or not; this is a judgment call. Furthermore it is a judgment call that will be shaped by a variety of unpredictable factors, such as how hard it is to retain subjects in the study, how dire the communication difficulties are, the time frame for the funding, and so forth. The PIs back in North America will never be in a position to dictate these standards in advance, or to know which standards were used in each case. Crucially, ensuring honesty, competence, and the free flow of information will not solve this problem.

In sum, since different researchers are working in different locations, using different disciplinary expertise, are driven by different working understandings of the goals of the research and subject to different local pressures resulting in varying standards, there seems to be no built-in guarantee that the justificatory stories that undergird the various pieces of the study together form one coherent justification, for which the group as a whole can be accountable. If the represented results are challenged, there may be *no* single justificatory story to be told about the methodological choices made and the epistemic standards used - not even one that would need to be told piecemeal by the various participants. Thus it seems like there is no reason to believe that the group collectively can be accountable for the finished product. Hence there seems to be little reason to count them as forming a group author.

Before moving on we will briefly consider a different way of bringing radically collaborative research under the 'centralized control' model. We develop this alternative and argue against it in much more detail in a different paper (Winsberg, Huebner, & Kukla 2014). Perhaps the central controller need not control each part of the research process and methodology, but instead the flow of information itself. That is, she might implement a social model that is designed to guarantee that the input from different researchers, operating according to different standards and in light of different pressures and interests, coheres into a reliable whole because of how it is combined. The reliability of the results would be defended not in terms of the specifics of the methodology of each part, but in terms of the reliability of the system for combining them – in something like the way that we can count on the 'wisdom of crowds' even (and indeed particularly) when we have no reason to count on the wisdom individual crowd members.

We cannot address this suggestion in detail here, but notice two things. First, as things currently stand, radically collaborative sciences like climate science and biomedicine have nothing resembling such a social model. They have not even acknowledged that there is an epistemic issue concerning how the different parts of such research projects combine. Second, whoever was in control of such a project would have to have an explicit commitment to her social model and a justification of its reliability. If we are interested in accountable authorship then it is not enough for us, from the outside, to decide that a result is reliable despite its inner chaos, perhaps for 'wisdom of crowds' sorts of reasons. Rather, the central controller would count as epistemically accountable only insofar as she was committed to the reliability of such a social model and ready to defend it or give it up as needed. In any case, if we think of *authors* as primarily authoring a system for coordinating parts whose individual reliability cannot be defended, then we are radically changing our conception of scientific authorship.

4. How deep is the 'no shared epistemic standards' problem?

In the preceding section, we talked about the problem of there being no shared epistemic standards that can be used to (more or less mechanically) assess the methodological choices of various collaborators. Broadly speaking, there are three things we could mean by "shared epistemic standards."

First, we might be making the familiar claim that there is no logic of induction. It is now virtually a truism in philosophy of science that there exists no single articulable set of rules that can deliver, mechanically, probabilities for a hypothesis given a specification of the evidence available in favor of (and against) it. This is a sense in which there are "no shared epistemic standards": if there rules for the epistemology of science, they are "one off." There might be, in other words, rules of thumb that can sanction empirical inferences in particular restricted domains, but these rules are brittle, and they do not travel well. They also, it is often argued, involve tacit knowledge that involves years of training to acquire.

This kind of lack of shared standards poses obvious difficulties for the sorts of strategies discussed in the last section, but they are not obviously insurmountable. Even if epistemic standards are fragmented and difficult to communicate, it might, in principle be possible for some centralized authority to master them. So the argument against easy answers to the author problem from this kind of failure of shared standards exists, but is relatively weak.

A second way of fleshing out the claim that there exist no shared epistemic standards would be to maintain that even given a particular epistemic task, a hypothesis, a body of background knowledge, and a body of evidence, there exists no uniquely correct answer concerning whether the evidence sufficiently supports the hypothesis. This version of the "no shared standards claim" is full-blown epistemic relativism. We reject this, or at the very least, we recognize that such relativism is incompatible with our notion of authorship, which centrally involves the concept of epistemic accountability. Genuine accountability involves being able to call the accountable person onto the mat and ask whether or not the relevant standards have been met; the very ideas of accountability and authorship makes no sense if there are literally no common standards at all. I cannot be held accountable for my choice of ice cream flavor.

But there is a third notion of "shared epistemic standards" that we think does fail to obtain, and whose failure poses a more serious and specific problem for radically collaborative research than the mere absence of a logic of induction. Note that even if there were a logic of induction, no scientific hypothesis would ever be established with certainty on the basis of a finite amount of evidence. The judgment whether a hypothesis should be accepted in light of the evidence involves two kinds of trade-offs. First, let us call the capacity of an investigation to generate persuasive results given a certain of effort and resources (and hence to collect a certain about of data), the "power" of the investigation. Methodological choices involve a trade-off between the power and the reliability of the investigation. Second, even given a choice along that continuum, methodological choices involve a further trade off between a desire to avoid false positive conclusion and a desire to avoid false negative conclusions – a so-called balance of inductive risks (Churchman 1948; Rudner 1953, Douglas 2000, 2004, 2009; Wilholt 2009) Decisions about how sure we need to be to accept a hypothesis always depend on the seriousness of making a mistake. Thus there cannot be any general or objectively correct answer to how such a trade-offs should be balanced. Values and interests inevitably play an important role in determining the 'seriousness' of accepting false positives and rejecting false negatives: to use

Richard Rudner's classic example, our concern with false negatives will be much higher if we are studying whether a drug has a potentially lethal side effect than if we are testing for defects in a set of belt buckles (Rudner 1953).

Since scientific inference is always uncertain, there is no avoiding making a judgment about how to balance inductive risks, and particular interests and investments will often affect this judgment: a drug company may well set a lower standard for 'establishing' efficacy than a consumer group. Furthermore, Douglas (2000, 2009) and others have argued, such interest-dependent judgments occur throughout the research process. How data are classified and coded, which sorts of screening tests are used, which methods are employed in smoothing and correcting data, and indefinitely many other judgments involve this type of inductive risk balancing. For example, whether researchers classify slides of rat tumors as benign or malignant turns out to depend, to a very large degree, on the goal of the study and their stake in the outcome (Douglas 2004).

In such situations, there is no value-free notion of a *correct choice* of distribution of inductive risks, since *any* choice will involve trade-offs. Likewise there is no value-free notion of "the correct epistemic standard," even in the "one-off" sense discussed above.[11] But remember where this leaves us: not only do different researchers in a radically collaborative research project use different methodological standards and make different choices, and not only do we *need* them to be able to exercise their judgment in this way, but there is no *right answer* as to how they should make these choices – no universalizable set of principles by which their performance can be measured. Each researcher faces micro- and macro-pressures that yields a local, one-off set of interests and values, and these will inevitably shape her choices all through the research process. Not only can we not typically recover these interests and choices (which, as we have pointed out, are likely often opaque to the researchers themselves) but even if we could, we would have no stable meta-standard for measuring their scientific reliability. Indeed, the relevant notion of reliability is at the center of what requires individual, value-laden judgment.

What sorts of interests and values might make a relevant difference to methodological choices, including inductive risk balancing, in the context of radically collaborative research? Climate science and biomedicine are normatively rich sciences that are driven by social values and demands in obvious ways. We have direct practical interests in the results of these sciences – interests that are personal, ethical, political, and economic. Different stakeholders obviously care in various ways whether a climate model gives

[11] This is not the same thing as full-blown epistemic relativism, since we can each still offer *reasons* for having chosen one methodological option over another.

encouraging or alarming results, whether a drug is more effective than the current standard of care, and so forth. A great deal of attention has been devoted to exploring how profit motives shape industry-funded clinical drug trials, and it may be possible to address the systematic structuring of studies by clear, recoverable interests (cf., Healy & Cattell 2003, Elliott 2011a, 2011b; Kahn 2004, 2006; Sankar & Kahn 2005; Sismondo 2007, 2008, 2009). But what matters more for our purposes is the micro-interests and micro-values that shape individual researchers' local practices, rather than the large-scale social investments that shape the research project as a whole and might be relatively easier to uncover and critically assess.

When bioethicists and philosophers of science have worried about the role of values and interests in shaping methodological choice points in biomedical research, they have focused overwhelmingly on *financial* interests. Indeed, concerns over the influence of financial pressures are justified. Private industry plays a huge and ever-increasing role in funding biomedical research. Scandals over ghostwriting, selective data use, selective publication, and ad hoc study designs are common, and often focus on the role of pharmaceutical companies and for—profit companies in controlling both the methods and the message of the science they fund. These scandals typically reveal a top-down, organized attempt to shape publications in ways that will increase the profits of a specific stakeholder. But published studies in biomedical science typically disclose dozens or even hundreds of conflicting financial interests. Collaborators often receive funding and support in the form of shares, grants, and lecturing fees, among many others. Collaborators from one part of a study are unlikely to know all of the micropressures that other participants are under to secure grants, satisfy donors, or wrap up a study.

Furthermore, financial interests are only one kind of pressure that shapes the methodological decisions and standards that are at play in biomedical research. At any given site, postdocs might be under pressure to please their advisers, there might be scientists who are attached to pet hypotheses or concerned to build their reputations, faculty might be under pressure to flesh out their annual reports or make tenure, and there may be disciplinary turf wars and grand ideological battles that guide the research. This is compounded by professional competitions and investments, along with pressures from local universities and communities. The course of the research is shaped by a wide variety of micro-interests, even where everyone behaves honestly and nothing is hidden. These interests are not typically coordinated with one another; often they may not be consciously available even to those that have them.

One source of such micro-interests – which bioethicists have discussed at length – derives from the fact that most collaborators in clinical research are healthcare professionals, so they come into the research project inhabiting dual roles with potentially conflicting duties. They are professional caregivers and healers, as well as scientists. But notice that these dual roles can create pressures that can affect inductive risk balancing and other methodological choices. For example, when faced with a subject who has a medical need, it is not obvious that the duty to stick closely to a scientific protocol will trump the duty to care for a person. In fact, Charles Lidz and his colleagues (2009) report that 64% of clinical researchers thought that researchers *should* depart from protocol to improve medical care, and many of them reported giving restricted medication to subjects, recruiting subjects they knew to be ineligible for a trial, and keeping subjects in a trial after they met the termination criteria in order to improve their health. As they note, such deviations may significantly affect the results of clinical trials.

5. Structured radical collaboration in high-energy physics

Up to here, we have focused on radically collaborative research in climate science and biomedicine. These are sciences in which it is not surprising that values and interests play a substantive role, given the socially pressing, normatively complex nature of their missions. Perhaps, then, our worries are not really about authorship in radical collaboration *per se*, but rather concern value-rich science. In this section we consider a 'pure' science that still operates through radical collaboration by necessity. Research in high-energy physics (hereafter HEP) must be distributed across numerous laboratories (which are often located in different locations throughout the world). In the case of the Organisation Européenne pour la Recherche Nucléaire (CERN), publications reporting collaborations on a single experiment may include as many as 1000 'authors', listed in alphabetical order without regard to seniority.[12] Like climate modeling and biomedical research, research in HEP relies on technological and symbolic resources that are rarely shared among different labs. The size and complexity of the detectors, the excessively long duration of the experiments

[12] The Stanford Linear Detector (SLD) worked to maintain "the coherence of the group and the *de facto* recognition that contributions to a piece of physics are made by all collaborators in different ways" (Galison 2003, 332). Each physicist was listed as a coauthor, and that the first publication on a particular topic included engineers as well. Authorship was not limited to those who 'contributed' to the reported result, but included everyone who had worked with a research group for a year (cf., Galison 2003, 334-335).

(some lasting as long 20 years), as well as the degree of collaboration required to produce a result, make isolated authorship impossible. But HEP is not clouded by policy entanglements, and indeed it seems to have little if any foreseeable practical import. Thus HEP is a good candidate for a science that may be radically collaborative without raising worries about the structuring role of multiple interests.

But consider an anecdote that was relayed to us. There were two major groups looking for the Higgs particle at CERN: ATLAS and CMS. When ATLAS reported data consistent with the observation of the Higgs particle at 2.8 sigma, the buzz around CMS was that they needed to do whatever was necessary to "boost their signal." This meant shifting their distribution of inductive risks to prevent them from falling too far behind the ATLAS group—toward higher power, at the expense of reliability, or towards a lower probability of missing a Higgs-type event, at the expense of a higher probability of finding a false positive. Hence even here, it seems that we see the influence of local pressures and interests on methodology, in ways that cannot be simply eliminated.

It may seem that such standards could be chosen collectively, or managed centrally, in ways that preserve the methodological transparency that is necessary for collective authorship. But a little bit of background should disabuse us of such a fantasy. The LHC detectors generate huge amounts of raw data: approximately 23 petabytes/second of raw data.[13] Hence, a triggering mechanism must be used to decide which data to keep, and which to ignore. Strikingly, this triggering process involves human judgment. When collisions occur, one cannot simply see a Higgs particle. At best, one sees the spray of particles that a Higgs event would produce if it occurred. But we are unable to calculate from first principles what this spray will look like, because we don't have a good enough theory of the strong interaction. So the calculation has to be semi-empirical, and some of the relevant empirical facts have to be observed on the fly. Very time-sensitive judgments have to be made, and they have to be made by widely distributed agents. And of course, every judgment involves a distribution of inductive risks. Thus it is hard to know how a perception that "we are falling behind the ATLAS group in the race" will affect the judgments on the ground.

This is something that people at CERN implicitly understand. They attempt to manage the role that interests might play in shaping perception by, for instance, using blinding procedures: The people making the empirical judgments about whether a Higgs event has occurred cannot see the part of the background

[13] Approximately 25 megabytes per event, times 23 events per beam crossing, times 40 million beam crossings per second in the center of the detector.

information that is relevant to the calculation of whether a “trigger” will be reached when they make their judgment. This separates the roles of those who decide how much to boost the signal from those who evaluate the resulting patterns; in other words, this structural mechanism is designed to minimize any distortion issuing from the interests of the scientists. Whether of not this technique helps to address this particular problem, it indicates a way in which inductive risk balancing continues to occur in unpredictable and perhaps unrecoverable ways throughout the research process, even where the research does not aim at some obviously value-laden goal.

Nonetheless, we think there are two relevant differences between HEP and our other two cases that make us worry less about the reliability of the results of HEP. 1) In HEP, the distribution of inductive risks is linear, in the sense that the Higgs Boson either exists or it does not. Compare to the climate case, where there are a wide variety of desirable prediction successes: global mean surface temperature, precipitation, sea level, ice melt, drought, storms, and other regional effect, etc. These are inductive risk considerations that pull in many directions; researchers can have any of a wide, multidimensional array of investments in various outcomes and parameters at both the micro and the macro level. In HEP, there is only one dimension of pull: you are either erring on the side of false positives or false negatives, when it comes to finding the Higgs Boson. This would seem to make the problem of ‘conflicting standards’ much more manageable, at least in principle. 2) Relatedly, in HEP, while there is a distribution of inductive risks at any given moment in time, in the long run, there will not be one. As more and more data are collected by the LHC, the power of the experiment gets larger and larger, and the balance of DIR gets smaller and smaller, until eventually, the choices will not matter. Likewise, the phenomena in HEP do not themselves change, so any attempts to understand them are, at least in principle, both cumulative and reproducible. In contrast, in climate science and biomedical research, the central questions are fundamentally time sensitive, both because the research questions are inherently local and the phenomena under scrutiny themselves dynamic. Of course, by 2100, we will know with certainty what the 2100 climate looks like. And in 2100, we will almost certainly retroactively know how effective some particular drug was, in some population. But our biomedical research practices are, of necessity, set up to give us answers *now*. HEP is set up to give us good results as soon as possible, but the data will continue to trickle in until we are, for all intents and purposes, sure of their implications.

While (at least in the long run) we may be less concerned about the *reliability* of HEP than of climate science or biomedicine, we have not found great reason to be comforted about the prospects for traditional authorship in this domain. Of necessity, a widely distributed, massive group of highly specialized researchers must participate in these projects. These researchers are not in a position to either understand all their colleagues' contributions or to track and vouch for their local choices and judgments, and it still seems unlikely that this problem will be eliminated by trying to make the science more mechanistic or transparent.[14] That we can be more confident that HEP will converge on reliable conclusions eventually is politically comforting, perhaps, but it does not help give us a model of authorship in radically collaborative research that should comfort the social epistemologist.

6. Conclusion

We have argued that there are—at present—no plausible strategies for understanding or ensuring accountable authorship in the context of radical collaborative research. We've tried to show that there is no one who is in a position to know, understand, or explain all of the epistemically significant decisions that have been made in the production of a purported piece of scientific knowledge. In radical collaborations, catch-and toss authorship cannot be sustained because the individuals who are involved in a research project cannot recover all of the epistemically significant decisions that have been made along the way; similarly, the necessity of local judgments that are made on the fly precludes the possibility of top-down, centralized control over knowledge production. Finally, as we suggested in the last section, the necessity of value-driven risk balancing decisions – which show up even in a 'pure' science like HEP – make genuine *group authorship* difficult to sustain. Together, these arguments suggest that neither individual nor multiple nor group authorship is plausible in radically collaborative research, at least in anything like its current organizational form. As things stand, authorship *simpliciter* remains impossible in most radical collaborations.

We have focused on the ways in which these issues arise in climate science and biomedical research. This is because the data that are collected and reported in these sciences are epistemically, morally, and politically important. Scientists working in these areas confront issues that are both socially pressing and normatively loaded, and it matters deeply that the data they reported are accurate and that there is someone or some people who can take responsibility

[14] Indeed, high-energy physicists have been exceptionally interested in authorship questions, and have developed various techniques for regulating and streamlining the authorship process; see footnote 11. But these have not been targeted at the kinds of concerns we have focused on here.

for this accuracy. This is not merely because we want to get the facts right, but because we often need to decide—right now—whether to remove a particular drug from the shelves, or whether to implement a policy that can sustain public health in the face of global climate change. The ability to interpret the results reported in radically collaborative publications, and to decide when and how to trust these results, matters not just epistemologically but also morally and politically.

Until now, radically collaborative science *as such* has not been identified as a distinctive phenomenon with its own epistemological and ethical challenges and complications. But such collaborations dominate many of the most socially pressing areas of science. This form of science is increasingly well represented in top journals and it is receiving more and more funding. We contend that the existing models of collective authorship and collaboration are unequipped to deal with the complexities that arise when epistemically robust values and judgments must be made at many disparate points in the production of radically collaborative data.

We also maintain that the models of collaboration and authorship that are applicable to small groups, where everyone can understand the judgments and decisions that are made by others, cannot simply be scaled up to address the issues that we have addressed. Importantly, this is not because we don't yet know how to scale then up. It is because scaling up is not the right way to address the issues we have raised here. These form novel class of epistemic and ethical challenges that arise specifically in sciences of scale. They are not, even in principle, the kinds of issues that we could scale up simpler models to address; they originate in the messiness that is produced by massively complex, distributed research architectures. Social epistemologists need to think creatively about authorship and accountability in the context of radically collaborative research, in which the production of a scientific result requires multiple, disparate kinds of experts, widely spread out across various locations, working under time- and context-specific conditions.

7. Works Cited

Andersen, H. & Wagenknecht, S. (2013). "Epistemic dependence in interdisciplinary groups," Synthese, 190: 1881-1898.

Biddle, Justin, and Eric Winsberg. (2009). Value Judgments and the Estimation of Uncertainty in Climate Modeling. In *New Waves in the Philosophy of Science*, ed. P. D. Magnus and Jacob Busch, 172–197. New York: Palgrave MacMillan.

Clark, A. (1998). *Being there: Putting brain, body, and world back together*. Cambridge: MIT Press.

Churchman, C. W. (1948). *Theory of Experimental Inference*. New York: Macmillan.

Douglas, H. (2000). Inductive Risk and Values in Science. *Philosophy of Science*, 67 (4), 559-579.

Douglas, H. (2004). Prediction, Explanation, and Dioxin Biochemistry: Science in Public Policy. *Foundations of Chemistry*, 6 (1), 49-63.

Douglas, H. (2009). *Science, Policy, and the Value-Free Ideal*. Pittsburgh: University of Pittsburgh Press.Dunne, J. (2006). Towards Earth System Modeling: Bringing GFDL to Life. Presented at the *ACCESS 2006 BMRC Workshop*. http://goo.gl/lPnwz (accessed 11 Jan. 2011).

Elliott, K. (2011a). *Is a Little Pollution Good for You? Incorporating Societal Values in Environmental Research.* (New York: Oxford University Pres).

Elliott, K. (2011b). "Direct and indirect roles for values in science," *Philosophy of Science*, 78: 303-324.

Galison, P. (2003). "The Collective Author," in *Scientific Authorship: Credit and Intellectual Property in Science*, Mario Biagioli Peter Galison, eds. (London: Routledge, 325-353).

Gilbert, M. (1989). *On Social Facts*, London: Routledge.

Gilbert, M. (1996). *Living Together: Rationality, Sociality, and Obligation*, Lanham, MD: Rowman and Littlefield.

Grier, D.A. (2005). *When computers were human.* Princeton: Princeton University Press.

Gowers, T. & Nielsen, M. (2009). "Massively collaborative mathematics," Nature, 461, 15: 879-881.

Healy, D. and D. Cattell. (2003). "Interface between authorship, industry and science in the domain of therapeutics," *The British Journal of Psychiatry*, 183: 22-27.

Huebner, B. (2013). *Macrocognition: Distributed Minds and Collective Intentionality*. (New York: Oxford University Press).

Kahn, J. (2004). "How a drug becomes 'ethnic': law, commerce, and the production of racial categories in medicine." *Yale Journal of Health Policy, Law, and Ethics,* 1: 1–46.

Kahn, J. (2006). "Race, pharmacogenomics, and marketing: putting BiDil in context," *The American Journal of Bioethics*, 6 (5): W1-W5.

Kukla, R. (2012). “Author TBD”: Radical collaboration in contemporary biomedical research,” *Philosophy of Science*, 79 (5): 845-858.
Lenhard, J. and E. Winsberg (2010). Holism, Entrenchment, and the Future of Climate Model Pluralism. *Studies in the History and Philosophy of Modern Physics* 41 (3), 253-62.
Lidz et al (2009), “Competing Commitments in Clinical Trials”, *IRB: Ethics and Human Research* 31:5, 1-6.
List, C. and P. Pettit. (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. (New York: Oxford University Press).
Poordad, F., et al. for the SPRINT-2 Investigators. (2011). “Boceprevir for Untreated Chronic HCV: A Randomized Clinical Trial.” New England Journal of Medicine 363:1195–1206.
Rolin, K. (2008). “Science as collective knowledge,” *Cognitive Systems Research*, 9 (1-2): 115-124.
Rudner, R. (1953). The Scientist Qua Scientist Makes Value Judgments. *Philosophy of Science* 20 (1), 1-6.
Sankar P. and J. Kahn (2005). “BiDil: Race Medicine or Race Marketing?” *Health Affairs*, W5 455-463
Shackley, S., J. Risbey, P. Stone, and B. Wynne (1999). Adjusting to Policy Expectations in Climate Change Modeling: An Interdisciplinary Study of Flux Adjustments in Coupled Atmosphere-Ocean General Circulation Models. *Cli* Shackley *matic Change* **43** (2), 413-454.
Shieber, J. (Forthcoming). Toward a Truly Social Epistemology: Babbage, the Division of Mental Labor, and the Possibility of Socially Distributed Warrant. *Philosophy & Phenomenological Research*.
Sismondo, S. (2007). “Ghost management: how much of the medical literature is shaped behind the scenes by the pharmaceutical industry?” *PLoS Medicine,* 4 (9): e286.
Sismondo, S. (2008). “How pharmaceutical industry funding affects trial outcomes: causal structures and responses,” *Social Science and Medicine*, 66 (9): 1909-1914.
Sismondo, S. (2009). "Ghosts in the machine: publication planning in the medical sciences,” *Social Studies of Science*, 39 (2): 171-198
Thagard, P. (2002). “Collaborative knowledge,” *Noûs*, 31 (2): 242-261.
Tuomela, R. (2002). *The Philosophy of Social Practices: A Collective Acceptance View*. (Cambridge: Cambridge University Press).
Weart, S. (2010). “The development of general circulation models of climate,” Studies In History and Philosophy of Science Part B, 41(3), 208-217.
Wilholt, T. (2009). Bias and Values in Scientific Research. *Studies in History and Philosophy of Science*, 40 (1), 92-101.

Winsberg, E. (2006). Models of Success vs. the Success of Models: Reliabilitywithout truth. Synthese, 152: 1-19.Winsberg, E. (2010). *Science in the Age of Computer Simulation*. Chicago: University of Chicago Press.

Winsberg, E. (2012). Values and Uncertainties in the Predictions of Global Climate Models. *Kennedy Institute of Ethics Journal*. 22 (2): 111-37.

Winsberg, E., B. Huebner, and R. Kukla. (2014). “Accountability, values, and social modeling in radically collaborative research.” *Studies In History and Philosophy of Science: Part A.* 46: 16-23.

Wray, K. (2002). “The epistemic significance of collaborative research,” *Philosophy of Science*, 69 (1): 150-168.

Wray, K. (2006). “Scientific authorship in the age of collaborative research,” *Studies in History and Philosophy of Science Part A*, 37 (3): 505-514.