

FORMAL SOCIAL EPISTEMOLOGY WORKSHOP

Cailin O'Connor, Assistant Professor

Logic and Philosophy of Science, UC Irvine

SCHEDULE OF TALKS

Friday, May 25th

8:30, Light Breakfast

9 am, Jeffrey Barrett (UCI), "Self-Assembling Networks"

(research with Brian Skyrms and Aydin Mohseni)

This talk brings together two new and developing research programs, the theory of self-assembling games and the theory of network epistemology. We will consider how an epistemic network might self-assemble from the ritualization of the individual decisions of simple heterogeneous agents. In such evolved social networks, inquirers can expect to be significantly more successful than they could be investigating nature on their own. We will consider networks that self-assemble in the context of both perfect and imperfect communication and compare the behavior of the modeled inquirers in each. We will pay particular attention to cases where a self-assembled network might be expected to be optimal.

10:15, Carole Lee (UW), "Gender-based Homophily in Scientific Collaborations"

Collaborators: Samuel Wang, Jevin West, Carl Bergstrom, Elena Erosheva

Scholarly collaboration now drives the frontier of knowledge. Given increasing professional advantages of collaboration such as those related to innovation, it is important to understand gender gaps in women's participation in collaborative teams. Using a rich database of publications in the JSTOR corpus, we explore gender-based homophily -- the tendency for researchers to co-author papers with individuals of the same gender -- in scholarly collaborations. Previous empirical efforts were limited to single academic disciplines. We make use of hierarchical representations of citation networks to investigate gender collaboration patterns across the scholarly landscape and at all levels of science, from narrowly focused communities to disciplines. We find that most disciplines contain a number of micro-communities with substantial behavioral gender-based homophily in co-authorship. The significant homophily at the micro-level propagates up the hierarchy. Focusing on the discipline of economics, we confirm previous findings that gender-based homophily increases with increased numbers of women in the subfield.

11:30, Liam K. Bright (CMU), "Specialisation"

In order to carry out their research scientists typically go through long and intensive training or educational programmes. This happens at grad school, but there are various senses in which continuing education is a normal part of scientific life. Think, for instance, of taking a sabbatical to learn a new technique or immerse oneself in a new literature. This is the process by which scientists continually specialise, and how scientists decide which specialisms to focus upon determines what skills, knowledge bases, and forms of expertise shall be available in the immediate future of a research field. In this paper my coauthor, Kevin Zollman, and I discuss the social-epistemic significance of these decisions. We produce and analyse a model of some aspects of these choices of specialisation. Within our model one finds what we call "the bushel effect", wherein under many configurations scientists end up systematically hiding their light under a bushel, which is to say not developing their talents to their full capacities and thereby imposing an epistemic opportunity cost on the epistemic community. We end by suggesting further research that might help uncover the extent to which actual processes of specialisation exhibit a bushel effect.

12:30-2:00, Lunch Break

2:00, Mystery Speaker, Title /Abstract tba

3:00, Coffee Break

3:15, Olufemi Taiwo (UCLA), "Discourse Power", (Liam Kofi Bright and Olúfẹ́mi O. Táíwò)

Communicative acts have been modeled in analytic philosophy of language as the exchange and group evaluation of propositions. The uptake such acts receive is affected by the social context in which communication occurs. Analytic philosophers of language, postcolonial theorists, and feminist theorists have all entertained the possibility that power hierarchies between the speakers, when salient, can produce predictable effects on communicative acts; including the ability of socially powerful persons to "silence" or distort the communicative acts of their interlocutors. In this paper, we produce an analysis of the mechanisms that might underlie such effects, which we term "discourse power". We show that a simplified model of communication that includes discourse power retains all of the predictive power of an analogous approach based on Stalnaker's notion of "common ground", but can also predict the silencing and distorting effects investigated by the aforementioned groups of theorists. Further, our model suggests there are as yet under-explored connections between discourse power in communication and agenda setting in a voting context.

3:45, Darcy McCusker (UW), "The Impact of Gendered Citation Practices"

Citations and other measures of publication impact are becoming increasingly important in decisions about science, such as employment, tenure and grant applications, but have not been the focus of much philosophical attention. One reason that citations may be worthy of philosophical attention is that citation metrics in a number of scientific fields show disparities by gender: women are cited less frequently than men, even when controlling for the number of women in a discipline. While the causes of such disparities are difficult to determine, the impacts on women in science are real. Establishing that even relatively small disparities in citations rates can have substantial impacts across the career of particular women or to the group of women in a discipline may be possible with formal tools. Understanding the potential impacts of gendered citation practices is especially useful in trying to determine the nature of the harm done to individual women, the group of women in a discipline, and to the discipline itself.

4:15, Eliana Peck (Vanderbilt), "Shame and Antiracism: On the Affective Features of Resisting Active White Ignorance"

Many thinkers have discussed the affective and epistemic traits that accompany American white people's investments in white supremacy; more recently, some have examined whether certain emotions might play a role in antiracist efforts by white people. This paper considers the contested emotion of shame, and describes areas of weakness in some of the existing literature. Drawing on newer research, the view argues for a limited, but positive, role for shame in the process of resisting active white ignorance, a form of produced ignorance that allows many white people to be complicit in racist harms while avoiding the negative affect that, for morally conscientious people, might otherwise accompany wrongdoing.

4:45, Aydin Mohseni (UCI), "Truth and Conformity on Networks"

Typically, public discussions of questions of social import exhibit two important properties: (1) they are influenced by conformity bias, and (2) the influence of conformity is expressed via social networks. We examine how social learning on networks proceeds under the influence of conformity bias. In our model, heterogeneous agents express public opinions where those expressions are driven by the competing priorities of accuracy and of conformity to one's peers. Agents learn, by Bayesian conditionalization, from private evidence from nature, and from the public declarations of other agents. Our key findings are that

networks that produce configurations of social relationships that sustain a diversity of opinions empower honest communication and reliable acquisition of true beliefs, and that the networks that do this best turn out to be those which are both less centralized and less connected.

Saturday, May 26th

8:30, Light Breakfast

9:00, Remco Heesen (Cambridge), “Why Should Scientists Do Good Science?”

The recent reproducibility crisis has revealed that a lot of research produced by scientists is not as reliable as we would like. This has led to various calls for scientists to do “better” science, e.g., sample size and power should be increased, false positive rates lowered, etc. As social epistemologists have noted, granting that such practices would make better science does not entail that it is better for scientists to engage in them. We are faced with the question: under what circumstances is it rational for scientists to do good science?

Two answers have traditionally been given. First, if scientists are motivated by purely epistemic considerations then it will be rational for them to engage in practices that are likely to be epistemically successful, i.e., do good science. Second, if scientists are motivated by a desire for fame then it will be rational for them to do what is most likely to get them fame, i.e., do good science. I use a combination of case studies and formal models to cast doubt on both answers: truth-motivated and credit-motivated scientists alike may be rational in doing bad science. I conclude that, even if we all agreed what practices constitute good science, we currently do not sufficiently understand how to motivate scientists to engage in them.

10:15, Manuela Fernández Pinto (Universidad de los Andes), “Should We Worry About Industrial Selection?”

The growing privatization of scientific research has raised important concerns for philosophers of science, especially regarding conflicts of interests and scientific integrity. Recently, however, Holman & Bruner (2017) have shown that an industrial selection bias can emerge in a scientific community without corrupting any individual scientist. The phenomenon is troublesome, for it seems to extract responsibility from scientific agents, relocating it in industrial decision making. In this paper, we examine the strength of industrial selection using a reinforcement Q-learning model, which simulates the process of industrial decision making when allocating funding to scientific projects. We show that, when there is no *a priori* knowledge about researcher bias, the industry is not so successful at allocating funding to industry-biased scientists due to the additional difficulty of inferring these biases from on-going scientific results. This in turn makes the industrial selection bias much weaker than previously conceived. If the result is correct, we can re-open the door to considerations regarding conflicts of interests and the individual responsibility of scientists.

11:30, Katharine Anderson (CMU), “Introducing Complexity to Economic Models of Collaboration and Knowledge Production”

Economists use the term “human capital” to represent everything that an individual brings to a productive process: their skills, abilities, experience, and knowledge. Traditional notions of human capital are born out of manufacturing, where we use single dimensional measures of a person’s speed, education, or tenure. Although simplistic, this way of thinking has generated considerable insight into behavior and outcomes in the production of things. However, the model starts to break down when it comes to the production of new ideas, where individual skills are often used in combination, shifting the attention from skill levels to skill sets.

Considering skills in combination makes the relationship between human capital and production much more complex: some skills are useful in combination, while others are not; some skills are common individually, but rare in combination. This means that a person’s contribution to knowledge production may be greater than the sum of

their individual skills, and the value of an additional skill might depend on the skills they already have. Here, I introduce some of this complexity into economic models of collaboration and knowledge production. I use a theoretical model of skill recombination to show that even when skills are assigned at random, the relationship between skill sets and outcomes is highly non-linear: small differences in skill sets are exaggerated to create large differences in outcomes. In academic collaboration, this creates superstars. In the labor force, it generates dramatic wage inequality.

Unfortunately, there is a considerable barrier to incorporating skill interactions into empirical models: even a small number of skills creates an intractably large number of interactions. I introduce a method of summarizing the mixture of skills in a population, using a network: nodes in the network are skills, which are connected by links if they are found in combination. This network can be used to create empirical measures of skill specialization, diversity, and coherence—measures of human capital that cannot be captured by existing models. These measures reduce the dimensionality of the problem, and allow us to consider more complex interactions between skills. I then illustrate these ideas using empirical examples: one collaborative, and the other from the labor market.

12:30-2:00, Lunch Break

2:00, Kevin Zollman (CMU), "Epistemic Governance"

In economics, there is a long standing discussion about market failure. Cases where an unregulated market will fail to achieve good outcomes for the participants are well known, and there is substantial debate about how to handle these situations. In this talk, I will show that analogous problems arise in epistemic contexts. Groups of inquirers (like scientists) can each attempt to learn the truth, but collectively fail to do so optimally. I will discuss how these problems arise and what we might do about them.

3:00, Coffee Break

3:15, Conor Mayo Wilson (UW), "Intersubjectivity and Statistical Principles"

Statisticians and formal epistemologists have defended a variety of "statistical principles" that characterize (i) when two data sets are "evidentially" equivalent and (ii) when available data "favors" one hypothesis over another. Such principles include the likelihood principle, the conditionality principle, the sufficiency principle, the law of likelihood, the stopping rule principle, and more. The number and variety of proposed statistical principles raises three questions: (1) How can one justify a purported evidential principle? (2) Are there more "statistical principles" that we have yet to discover? and (3) Do existing statistical principles have qualitative analogs? In this talk, I propose a methodology for answering all three questions. Statistical principles, I argue, ought to be justifiable as *agreements* that could be reached by hypothetical scientists with differing beliefs but who all satisfy basic coherence axioms. When the coherence axioms are the standard Bayesian ones, for example, then scientists will endorse the claim that the likelihood principle characterizes *precisely* when two data sets are evidentially equivalent, thus showing that further statistical principles (characterizing equivalence) are unnecessary. I conclude by describing a project: to find qualitative analogs of existing statistical principles by weakening the coherence axioms that members of one's hypothetical scientific community ought to satisfy.

4:15, Mike Schneider (UCI), "Priority and Privilege in Scientific Discovery" (Non-presenting coauthor: Hannah Rubin)

The priority rule in science has been interpreted as a behavior regulator for the scientific community, which benefits society by adequately structuring the distribution of intellectual labor across pre-existing research programs. But considerations about how news of scientific developments spreads throughout a scientific community at large suggest that the priority rule is something else entirely, which disadvantages historically underrepresented or otherwise marginalized social groups.

4:45, Rachel Cooper (UCI), "Thinking About Social Mechanisms: Can Social Systems be Mechanistically Explained?"

My aim in this paper is to make a case for thinking that social systems can be mechanistically explained. The particular phenomena of interest are regularities among members of human kinds, such as semi-uniform behaviors and positions of privilege and disadvantage. I consider my goals to be in line with those who have called themselves "Debunkers": I want to highlight the lack of justification for the disadvantaged positions of members of certain groups, by showing that the supposedly natural features they share do not themselves explain such behaviors or positions, that these regularities are not inevitable. The motivating idea behind my approach is that, if the shared features do not themselves explain regularities among members, then social mechanisms might.

In order to motivate the project, I first place myself within the Social Construction literature, by giving reasons for my focus on social explanations as opposed to the individuation conditions for social categories, or the metaphysics of such categories. I next consider a few of the major developments in theories of Scientific Explanation within the last century, and wherein "social explanations" might fit in best. Because I am ultimately interested in making sense of the idea that there could be genuinely social causes, I focus on mechanistic explanation. I then extract some important concepts found within Social Ontology/Epistemology and place them within the "New Mechanisms" framework. I argue that this approach of breaking down the social systems of interest into entities and activities is beneficial in both explaining the regularities we find among group members, and providing insight into how a mechanism can fail to be inevitable: the active entities are social groups, which are comprised by agents, who could in theory learn to do otherwise.

I conclude that social systems can be mechanistically explained: the New Mechanisms framework allows us to understand how there could be genuine social causes, and yet also that the regularities produced by these mechanisms are—at least in principle—non-inevitable, in a way that those produced by other complex systems (e.g., biological ones) are not, in being broken down into entities that are ultimately comprised of agents, who—again in principle—could learn to do things differently. The upshot is that, in fitting the claims of Social Construction into a framework recognized by Philosophers of Science, we not only pave the way to wider acceptance of the legitimacy of such claims, we also are able to make clear the intuitively important difference between social and other complex systems: that we actively maintain them.

5:15, Kierstan Thomas (Penn State), "In Defense of Intersectionality as a Decolonial Method"

This paper serves as a critical engagement with María Lugones' essay "Radical Multiculturalism and Women of Color Feminisms" in order to uncover the decolonial lens within Kimberlé Crenshaw's theory of intersectionality. In her assertion that intersectionality is a "white bourgeois feminism colluding with the oppression of Women of Color", Lugones renounces intersectionality's alleged reliance on logics of oppression. She ultimately concludes that intersectionality fails an appropriate means of both resisting and deconstructing the colonial ideologies behind gender. I explore these claims by arguing that Lugones' conclusion relies upon the hidden claim that subjects cannot resist within structures of oppression, and must move outside of those structures before engaging in meaningful resistance. I then juxtapose Lugones' perception of intersectionality with Crenshaw's own claims, while also utilizing José Medina's concept of "epistemic resistance" to demonstrate how resistance operates in oppressive structures such as the colonial gender system.