

| | |
|---|----|
| Amin Ebrahimi Afrouzi—The Ontological Subjectivity And Objectivity Of State Law | 8 |
| Paul Dragos Aligica—Alternative Communities, Societies and Institutions..... | 9 |
| Facundo Alonso—Shared Intention, Morality, and Relational Normativity | 11 |
| Line Edslev Andersen and K. Brad Wray—Collective Authorship: What We Can Learn From Retraction Notices | 12 |
| Derek Anderson—Gender Authenticity And Anti-Essentialism..... | 13 |
| Anne-Sofie Munk Autzen—Collaborative Improvisation and Shared Intention | 14 |
| Matthew Baddorf—A Theory of Collective Virtue | 15 |
| Paweł Banaś and Krzysztof Poslajko—There Are Only Two Kinds of Social Kinds | 16 |
| Donald Baxter—Social Complexes, Aspects, and Many-One Identity | 17 |
| Saba Bazargan-Forward—Collective Responsibility and Agentive Functions | 18 |
| Rose Bell—Inside and Out: Ascriptivism About Gendered Traits | 19 |
| Cristina Bernabéu—The Active Role of Technological Artifacts In Enacting Gender Norms | 20 |
| Sara Bernstein—The Metaphysics of Intersectionality | 22 |
| Sam Berstler—Social Acts and the Limits of Pretense..... | 23 |
| Matteo Bianchin—Ideology and Social Structures | 24 |
| Cristina Bicchieri—The Dynamics of Norm Erosion..... | 25 |
| Gunnar Björnsson—The Importance of Shared Obligations | 26 |
| Kurt Blankschaen—Are Mass Shooters A Social Kind?..... | 27 |
| Olle Blomberg—Responsibility Gaps and Attributionism | 28 |
| Sarah Braasch—Solving Lewis And Hart’s Problems with Normativity, Evolution, And Disagreement | 29 |
| Michael Bratman—Institutional Functioning, Institutional Agency..... | 30 |
| Thomas Brouwer—Social Contradictions | 31 |
| Åsa Burman—Telic Power..... | 32 |
| César Cabezas—The Role of Motivation in Structural Explanations of Oppression | 33 |
| John Cantwell and Tobey Scharding—Hegelian Collective Rationality, Recognition, and International Business..... | 34 |
| Elizabeth Cantalamessa—Disability Studies, Conceptual Engineering, and Metalinguistic Negotiation..... | 35 |
| Will Zhang Chen—Group Ontology and the Right Group Epistemology..... | 36 |
| Panagiotis Christias—The Social Ontology of Panajotis Kondylis: Metaphysics of The Social Being..... | 37 |
| Daniel Cole—Persistent Problems In (Some) Social Scientific Conceptions of Property And Law | 38 |

| | |
|--|----|
| Stephanie Collins—Blameworthiness for Non-Agent Groups..... | 39 |
| Rachel K. Cooper—On the Distinction Between Natural and Social Kinds | 40 |
| Sean Cordell—Institutions, Functions, and Ethics..... | 41 |
| Kevin Craven—Amelioration and the Ethics of Gender Ascription | 42 |
| Ben Curtis—Plural Subjects and the Question of Class Consciousness..... | 43 |
| Giuseppe Danese—Where Social Ontology Meets Economic Anthropology..... | 44 |
| Sven Delarivière—Groups As (Epistemic) Agents: A Conceptual Defence for Changing the Subject..... | 45 |
| Robin Dembroff—Social Categories and the Politics of Recognition | 46 |
| Alexander Dietz—Unit-Of-Agency Dilemmas | 47 |
| Zili Dong—The Reality of Institutional Kinds: A Causal Account..... | 48 |
| Hein Duijf—Responsibility in Equilibrium | 49 |
| Michelle Dyke—Rethinking What It Takes To Be An Agent..... | 50 |
| Adam Dyrda and Paweł Banaś—Were American Legal Realists Naïve Metaphysicians? | 51 |
| Douglas Edwards—Truth and Social Ontology | 52 |
| Brian Epstein—Convention and the Nature of Institutions | 53 |
| August Faller—How Do Social Groups Persist? | 54 |
| Jose Luis Fernandez—Circularity in Searle’s Social Ontology: With a Hegelian Reply | 55 |
| Roberta Ferrario and Daniele Porello—An Ontological Analysis of Groups as Complex Variable Embodiments..... | 56 |
| Peter Finocchiaro—Intrinsic Masking and Sexual Orientation | 57 |
| Christina Friedlaender—Ideal Theory and Shared Intentions | 58 |
| Caitlin Fitchett—Shared Intentions and Intending to Merely Do One's Task..... | 59 |
| Brian Flanagan—Intuitions About Group Minds: Public Evidence | 60 |
| Jade Fletcher—Can Gender Be Part of the Plan? | 61 |
| Carmel Forde—According to Others | 62 |
| Megan Fritts—An Argument Against Causal Explanations of Akrasia | 63 |
| Christopher Fruge—Against the Collective Responsibility of Groups..... | 64 |
| Barbara Fultner—Joint Intentionality and Embodiment: A Performative Account of Collective Action | 65 |
| Roberto Fumagalli—Rationality and the Acquisition of Intentions: Why Rational Choice Theory Is Not Self-Defeating..... | 66 |
| Jorge Garcia—Analyzing Social Construction, Dismantling Wide Constructionism | 67 |
| Benjamin Genta—On the Possibility of Group Forgiveness | 68 |

| | |
|--|-----|
| Samantha Godwin—Deconstruction of Gender, Reification Of Race: Inconsistent Theories of Social Ontology on the Left | 69 |
| Matthew Goldstein—Harmonic Ontologies | 70 |
| David Gray—Racial Objects: A Fanonian Account of Race | 71 |
| Aaron Griffith—Anchoring Human Social Kinds | 72 |
| Francesco Guala—Rescuing Individualism, or How Supervenience May Survive the Materialist Challenge | 73 |
| Mattias Gunnemyr—Political Responsibility Without Participation and Vice Versa | 75 |
| Niels de Haan—A Taxonomy of Prospective and Retrospective Moral Responsibility in Collective Contexts | 76 |
| Jani Hakkarainen—Constitution of Social Kinds As Generic Identity | 77 |
| Raul Hakli—Team Reasoning and the Irreducibility of Collective Intentionality | 78 |
| Mary Elizabeth Halper—A Platonic Perspective in Husserl’s Social Ontology | 79 |
| Jasper Heaton—Structuralism and Individualism Functioning Together | 80 |
| Marcus Hedahl—Writing Together: What the President’s Daily Brief Can Reveal About the Significance of Collective Written Testimony | 82 |
| Cassie Herbert—Derogatives: Beyond Insults and Slurs | 84 |
| Kendy Hess—Re-Bunking Corporate Agency | 85 |
| Frank Hindriks—The Duty to Join Forces | 86 |
| Frank Hindriks—Socially Constructed Kinds | 87 |
| Onni Hirvonen—Multi-Dimensional Personhood and Attributing Responsibility to Group Agents | 88 |
| Säde Hormio—Non-Agential Collectives and Consumer Complicity for Climate Change | 89 |
| Graham Hubbs—Individualism, Psychologism, and Collective Intentionality | 90 |
| Polycarp Ikuenobe—The Social Ontology of African Communalism and Personhood | 91 |
| Bernard Jackson—Awareness, Dependency, and the Heart of Blame | 92 |
| Genevieve Jacot-Guillarmod—Arguing for the Possibility of Group Self-Deception | 93 |
| Daniel James—Power as a Social Ability—Towards a Philosophical Explication | 94 |
| Bartosz Janik—Legal Positivism and Varieties of Dependence | 95 |
| Ludger Jansen—The Formal and the Informal. Towards an Integrative Theory of the Social | 97 |
| Katharine Jenkins—Understanding Oppressive Ontologies | 98 |
| Todd Jones—Culture Does Not Explain Behavior | 99 |
| Laura Kane—The Ontological Identity of the Family: Sociality, Purpose, and Care | 100 |
| Jeffrey Kaplan—How to Make a Rule: A Defense of a Supposedly Refuted Theory | 101 |

| | |
|---|-----|
| Jessica Keiser—Perduring Languages..... | 102 |
| Jo-Jo Koo—What Is the Relevance of Social Ontology for Social Scientific Research? | 103 |
| Odin Kroeger—The Metaphysics of Intersectionality | 105 |
| Nick Kroll and Bennett Helm—Discoverability and Contestability in Social Construction (Of Gender) | 106 |
| Visa Kurki, Säde Hormio, and Pekka Mäkelä—Collective Agents and Corporate Responsibility | 107 |
| René Lacroix—Is Face an Institution? | 108 |
| Lauri Lahikainen—Climate Change and the Collective Responsibility for Climate Change | 109 |
| Arto Laitinen—Social Wrongs..... | 110 |
| Arto Laitinen and Joanna Bryson—Collective Responsibility to Prevent AI Systems from Becoming Responsible | 111 |
| Dimitri Landa—Two Paths to We (As I) | 112 |
| Helen Lauer—Scientific Consensus and the Discursive Dilemma | 114 |
| Holly Lawford-Smith—Collective Responsibility Revisited..... | 115 |
| Edward A. Lenzo—Collective Moral Luck | 116 |
| Felipe León—Mutual Responsiveness in Action and Shared Action Spaces | 117 |
| Peter Levine—What Should We Do? The Bloomington School and the Citizen's Question | 118 |
| Ronald Loeffler—Belief and Linguistic Understanding | 119 |
| Kirk Ludwig and Marija Jankovic—Conventions and Status Functions | 120 |
| Taylor Madigan—Shared Agency and Mutual Accountability | 121 |
| Pekka Mäkelä and Raul Hakli—Hybrid Agents and Responsibility Sharing | 122 |
| Teresa Marques—Dangerous Speech and Collective Action | 123 |
| Annette Martin—Biased Minds as Internalized Structural Phenomena: A Dual Account of Social Structure | 124 |
| Laura Martin—The Social Ontology of Systemic Injustice | 125 |
| Eduardo Martinez—Populism, Social Groups, and Democratic Health..... | 126 |
| Russ McBride—Deontic Binding..... | 127 |
| Molly McGrath—The Intentional Structure of the Modern State..... | 128 |
| Torsten Menge—What Determines the Boundaries of a Political Community? | 129 |
| Kristina Meshelski—The Consequences of the Social Construction of Race for Social Contract Theory | 130 |
| Imko Meyenburg and Ana Turcitu—Vagueness and Social Ontology..... | 131 |
| Seumas Miller—Judgment, Agency and Joint Epistemic Action | 132 |
| Anna Moltchanova—Intentional Institutional Agency in Oppressive Societies | 133 |

| | |
|---|-----|
| Jonah Nagashima—What Counts as Fundamental? | 135 |
| Bernhard Nickel—The Division of Ideological Labor | 136 |
| Patrick O'Donnell—White Supremacy as Metaphysics..... | 137 |
| John Oduor—The Religious Conception of Social Ontology..... | 138 |
| Noriaki Okamoto—Interpreting Collectively Accepted Social Norms from a Perspective of Performativity..... | 139 |
| Alex Ozar—"Some Are Guilty, All Are Responsible" | 141 |
| Rohit Parikh—To What Extent Is a Group an Individual? | 142 |
| Clarisse Paron—In Defence of a Metaphysics of Race..... | 143 |
| Asya Passinsky—A Normative Hylomorphic Theory of Social Objects | 144 |
| Grace Paterson—Speaking for Another | 146 |
| Armando Perez-Gea—The Nature of the Political Association..... | 147 |
| Laura Perez—Visual Perception, Social Institutions and Corruption | 148 |
| Björn Petersson—We-Intentions and Immunity to Error Through Self-Misidentification | 149 |
| Lien Pham—A Case for Sustainable Development in Vietnam's Emerging Civil Society | 150 |
| Orlando Poblete—A Journey into the Epistemology of the Act of Judgement..... | 151 |
| Darien Polock—Beyond Pluralism: Racial Concepts, Pejorative Content, and Racial Eliminativism | 152 |
| Matthew Rachar—How We Act Together | 153 |
| Mike Raven—A Puzzle for Social Essences | 154 |
| Rene Reich-Graefe—Bayesian Trust, Trust Intermediation & Boundary Spanning | 155 |
| Jared Riggs—Saying What Race and Gender Are | 156 |
| Katherine Ritchie—Does Identity Politics Reinforce Oppression? An Answer from Social Metaphysics | 157 |
| Claudia Rosciglione—The Constitutive Aporia of the Relationship Between Individuals and Collectivity..... | 158 |
| Don Ross and Wynn Stirling—Modeling Group Agency and Social Norms Using Conditional Game Theory | 160 |
| Abe Roth—Directed Duties, Shared Agency, and Wronging Another | 161 |
| Joshua Rust—Institutional Identity | 162 |
| Mikko Salmela and Michiru Nagatsu—Interpersonally Scaffolded Affectivity: The Social Motivation Approach | 163 |
| Jules Salomone—Cooperation and Shared Intentional Actions | 164 |
| Kirun Sankaran—Adam Smith On Moral Sentiments in Structural Contexts | 165 |
| Gloria Sansò—Felicity Conditions of Document Acts | 166 |

| | |
|---|-----|
| Heidi Savage—The Metaphysics and Politics of Personhood: Issues in the Social Ontology of Persons | 167 |
| Cosim Sayid—Communication, Inc. | 168 |
| Naomi Scheman—From the Social World to Stones and Back Again..... | 169 |
| Hans Bernhard Schmid—The Sources of ‘We-Intention’: Sellars, Durkheim, and Collingwood | 170 |
| Kate Schmidt—Inclusivity as a Virtue of Social Groups..... | 172 |
| Michael Schmitz—Toward A Layered Account of Collective Intentionality..... | 173 |
| David Schweikard—Responsibility for Collective Inaction..... | 174 |
| Anne Schwenkenbecher—Collective Inaction and Group-Based Ignorance..... | 175 |
| Anne Schwenkenbecher—Do Group Agents Resemble Psychopaths? | 176 |
| Keyvan Shafiei—How Is This Stuff Fundamental?..... | 177 |
| Kenneth Silver—Group Action Without Group Minds | 178 |
| Yannick Slade-Caffarel—The Cambridge View of Emergence | 179 |
| Matthew Smith—Reliance as a Source of Shared Action | 180 |
| Sabine Spangenberg—The Reframing of Social Choice Theory Under the Consideration of Conditional Self-Interest..... | 181 |
| Todd Stambaugh and Rohit Parikh—Knowledge, Behavior, and Rationalizability..... | 182 |
| Cynthia Stark—Gaslighting: Individual and Collective, Epistemic and Manipulative | 183 |
| Daniel Story—Joint Action Without Theory of Mind | 184 |
| Megan Stotts—Toward A Behavioral Account of Social Institutions | 185 |
| Anna Strasser—Asymmetric Cases of Joint Actions..... | 186 |
| David Strohmaier—Group Agency and Homuncularism | 187 |
| Richard Strong—Clarifying the Role of Socialization and Cognition for Social Ontology... .. | 189 |
| Ásta Sveinsdottir—Social Construction, and the Problem of Passing..... | 191 |
| Larisa Svirsky—Responsibility as a Relational Property..... | 192 |
| James Swindler—Why Do We Care What Others Think? Should We?..... | 193 |
| Andras Szigetı—Judgment Aggregation and Collective Responsibility..... | 194 |
| Elanor Taylor—Naturalness, Context and Social Categories | 195 |
| Deborah Tollefsen—Group Knowledge Is Group Know How | 196 |
| Leo Townsend and Dina Townsend—Silencing Indigenous Groups..... | 197 |
| William Tuckwell and Paul-Mikhail Catapang-Podosky—Can Social Groups Bear Obligations To Combat Structural Oppression?..... | 198 |
| Raimo Tuomela—Constitutive Rules Are Not Reducible to Regulative Ones | 199 |
| Tuomas Vesterinen—The Explanatory Power of Natural Kinds | 200 |

| | |
|--|-----|
| Lovey Vikram—Understanding Collective Emotions | 201 |
| Jaana Virta—Social Construction of Gender in the Theory of Gender Performativity | 202 |
| Konrad Werner and Krystian Wiciarz—Towards an Ontology of the Echo Chamber | 203 |
| Randall Westgren—Identity, Identification, and Intention in Formal Organizations | 204 |
| Charlotte Witt—Voluntarism About Gender | 205 |
| Bill Wringer—What Can We Do? Collective Ability, Collective Obligation and Non-Agent Groups | 206 |
| Daniel Brinkerhoff Young—“Racism Without Races”—Socio-Cultural Racialization, Latinxs, and Muslims..... | 207 |
| Vojtěch Zachník—Why Conventions Cannot Be Social Institutions? | 208 |
| Aleardo Zanghellini—Conceptual Analysis and the Nature of Law | 209 |
| David Zoller—Which Room Am I In? Arendt and Husserl On Collective Contexts and Homeless Minds | 210 |

The Ontological Subjectivity and Objectivity of State Law

*Amin Ebrahimi Afrouzi—The Ontological Subjectivity And Objectivity Of State Law
University of California, Berkeley*

Many theorists have noted that ‘what law is’ depends, at least in one respect, on what is “adopted and accepted” in a given community as law. They go on to describe law variably as a function of habits of behavior (Bentham and Austin), social practices (Hart 1961, Dworkin 1986), social conventions (Postema 1982, 2012; Green 1999; Southwood 2011, 2017), and plans (Shapiro 2011). These different terms try to capture the feature of law that is socially constructed or “positive.” Much of this literature focuses on contrasting state law with moral (Hart 1961) and physical (Dworkin 1986) phenomena, which exist independently of intentional activity. As a product of intentionality, therefore, state law sits squarely at the center of social ontology.

However, it is not enough to say that state law is intentionally constructed, and I shall add that existing ontological phenomena, whether objective or subjective in their origins, both partly dictate appropriate intentional attitudes towards them and partly limit the possibilities of intentional construction. Thus, pace Dworkin and Shapiro, we must fully acknowledge that once intentionally constructed, ontologically subjective phenomena become congealed as a part of objective reality. Yet, pace mainstream theories of social ontology (i.e. Searle 1996), we must admit that ontologically subjective phenomena such as state law are limited both by the reality that exists independently of them (cf. Butler 1990, who denies this) and the reality that they constitute. This is important because 1) even if we could solve the problems of collective intentionality, equating state law’s ontology with intentional adoption and acceptance is a reduction at face value (cf. Searle 1996), and 2) in accounting for the normativity of state law, it would be too circular to say that our intentions are normatively bound by our intentions (cf. Butler 1990).

Following this observation, a tiered theory of law will start to emerge: the positive construction of state law is intentionally constructed in stages, wherein the existing ontological phenomena at every stage both partly dictates appropriate intentional attitudes towards them and partly limit the possibilities of further positive creation of those phenomena in that stage. In the remaining of the paper, I shall give content to the emerging skeletal sketch: exactly how state law limits or enables our intentionality, through what mechanisms can it be created and sustained, and to what extent these mechanisms are ontologically subjective.

Alternative Communities, Societies and Institutions: Social Ontology and the Cases of Deliberated Social Experiments In Creating Alternative Systems of Social Order

*Paul Dragos Aligica—Alternative Communities, Societies and Institutions
George Mason University*

What are the uses of Social Ontology -broadly defined- and of its conceptual apparatus and insights, when it comes to social and political analysis? What is the relevance of Social Ontology in comparison (and as value-added) to more traditional economic, social and political theorizing? This paper will be an attempt to explore these and related questions from a perspective coming from outside of the field of Social Ontology. It is an exploratory exercise of an outsider (who has a background in public choice institutionalism, political economy and political theory) to get familiarized, navigate and get a better sense of the relevance of the emerging field when it comes to social and political analysis.

The approach has two steps. First, in order to focus the perspective and to avoid discussing the problems on too broad and vague references to institutions, social reality and social practices, it will identify and add a new group of exemplary cases to the population of cases used in Social Ontology and Collective Intentionality research. The paper identifies one particular subset of massive relevance for the kind of problems, processes and approaches the Social Ontology agenda is interest in: The cases of deliberated social experiments in creating alternative systems of social order, alternative communities, intentional communities, alternative societies, alternative institutional arrangements etc.

The population of such “natural experiments” in intentional social construction, based on endeavors to rethink, redesign and reconstruct social order, is nontrivial and relatively well documented. We have thus access to a solid record reflecting, articulating and reporting the design, planning, implementation and assessment of those experiments, put to practice in diverse forms and diverse circumstances, at different scales and with various degrees of success. From large scale plans of recreating entire countries, nations and societies on Communist or Socialist lines, to the lower scale experimentation of Harmonists, Icarians, Aurora, Oneida and Wallingford perfectionists, the range of such designs and social experiments is remarkable. They are all offering -- both in their failure and success-- a window to the processes and practices involved in such exercises of collective intentionality, collective decision, collective action, and social construction and constitutive performances.

In addition to the simple idea of focusing the perspective and anchoring the discussion in real and concrete cases and examples, the strategy involved is based on the Nobel Prize in economics winner Elinor Ostrom’s approach to the use of case studies. Ostrom notes that we should learn from biologists’ approach when they are facing situations of studying complex processes that are poorly understood. The strategy is to identify cases “in which the process occurs in clarified even exaggerated forms”. The case “is not chosen because it is representative for all forms”. Rather the case is chosen “because particular processes can be studied more effectively using it than using another”.

The second step is to approach these cases through the conceptual lenses of Social Ontology. In this paper --again, with a view to make things more precise-- out of the broad and diverse Social Ontology literature, the focus will be mainly on one such specific set of lenses: Brian Epstein's Social Ontology apparatus, as articulated and presented in *The Ant Trap*. The underlying logic of the exercise is a reflection of the traditional social science double approach: Using a conceptual framework to illuminate, analyze, and interpret a case. While using the case as a vehicle to test and develop the conceptual framework.

In brief, the paper will try to put in dialogue a set of social phenomena of relevance for the basic themes of the Social Ontology literature, with a set of theoretical and conceptual lenses designed to analyze and increase our understanding of such phenomena.

Shared Intention, Morality, and Relational Normativity

Facundo Alonso—Shared Intention, Morality, and Relational Normativity
Miami University, Ohio

Philosophers have argued that what distinguishes joint action from other kinds of aggregated phenomena lies in its internal component, in the fact that in joint action participants have a joint or shared intention to act and that they act on that intention. It is a vexed question what the nature of shared intention is. In previous work I have argued that shared intention is both a psychological and a normative (moral) phenomenon (-----, -----). Although its spirit has been cashed out in different ways, the idea that shared intention is a psychological phenomenon has been widely accepted in the literature (-----, Bratman, Ludwig Searle, Tuomela). In contrast, the idea that shared intention is also a moral phenomenon has met considerable resistance (Bratman, Gilbert, Roth). In this paper I offer further reasons for thinking that shared intention is indeed (in part) a paradigmatic moral phenomenon. In the context of shared intention morality gets a grip on us at two levels. First, in sharing an intention to act together we interact with other people, who merit equal respect and concern. Secondly, and more importantly, the way in which we interact with others in shared intention is a basis for, or a source of, interpersonal obligations (-----, -----). I show that shared intention is like promising and friendship in these two respects. In addition, I defend the claim that the obligations created by shared intention are genuinely moral obligations. I do this in response to a familiar objection raised by Margaret Gilbert, to the effect that obligations of shared intention are “relational” (or “bipolar” or “directed”) and that this is incompatible with their being moral obligations. I argue, against Gilbert, that the obligations of shared intention are both moral in nature and relational in structure and, moreover, that it is precisely because the relational obligations of shared intention are of a moral nature that we can make proper sense of the type of special standing each participant in shared intention has to demand certain actions from the other participants and to complain to them if they fail to perform.

Collective Authorship: What We Can Learn from Retraction Notices

*Line Edslev Andersen and K. Brad Wray—Collective Authorship: What We Can Learn From Retraction Notices
Aarhus University*

When a serious error is discovered in a collaborative scientific paper and the paper has to be retracted, the group of authors – the collective author – is put under pressure. In this talk, we examine the robustness of collective authors by examining how they react to this kind of pressure and when they dissolve. This may help uncover the nature of collective authorship.

Our approach is to study retraction notices. We have looked at retraction notices in Science from the past 35 years and made a preliminary analysis of the notices signed by the authors of the retracted papers (as opposed to editorial retractions). They constitute 85 % of the retractions.

The preliminary analysis suggests that there are three common types of retractions. First, there is the type of retraction that is due to misconduct and where the collective author lays the blame on an individual who is guilty of misconduct [Category 1]. Second, there is the type of retraction that is due to an honest mistake and where the collective author takes on the responsibility, identifies the error made, and evaluates the difference the error makes to the reported findings [Category 2]. Finally, there is the type of retraction where it is unclear whether misconduct or an honest mistake has been committed and where the collective author neither obviously takes responsibility nor places it on someone else [Category 3].

If these are three common types of retractions, what do the various types tell us about collective authors?

Category 1 cases, where a collective author does not take responsibility when a group member has committed misconduct, may suggest a certain disunity of collective authors. But it may also suggest something else. The reason why the collective author does not take responsibility in these cases may be that a group member is not acting as a group member at all when committing misconduct, since her intentions conflict with the main aim of the group, assuming the main aim of the group is to uncover the facts about something.

Category 2 cases, where a collective author takes on the responsibility when an honest error has been made, suggest a certain robustness of collective authors. The collective author will describe the error it committed and reevaluate the reported findings.

Category 3 cases are likely due to the fact that a collective author will sometimes suspect that the error causing the retraction is due to misconduct. In this case, the collective author has reason not to trust what the group member who knows what happened says, since this is the person who is suspected of having committed misconduct, and thus of having lied to the group before. At the same time, the collective author will have moral and legal reasons not to publicly state mere suspicions. This makes for retraction notices where it is not clear who is responsible and why exactly the paper is being retracted.

Gender Authenticity and Anti-Essentialism

Derek Anderson—Gender Authenticity And Anti-Essentialism

Boston University

Feminist theory is nearly univocal in opposing gender essentialism. This opposition creates a *prima facie* difficulty for making sense of certain trans experiences: experiences of gender authenticity, as opposed to experiencing one's gender identity as inauthentic. Some trans people experience this shift to authentic experience when they begin to identify as the gender they in fact are. Whatever makes an experience of one's gender authentic as opposed to inauthentic seems to be some kind of gender essence, a substantive truth about what gender a person really is. Here I explore two aspects of the problem of essentialism for gender authenticity—anti-substantialism and intersectionality—and develop an account that navigates them.

Both issues arise in Judith Butler's *Gender Trouble*. First, substantialism—the idea that “man” and “woman” refer to ‘abiding substances’—is seen as politically harmful because it “subordinates dissonant gendered features.” The culturally constructed expectation that gender is a substantive aspect of a person gives rise to constraints on the performativity of gender, which in turn normalize the attribution of gender essence in accord with dominant cis-het ideology. Theories positing substantial truths about gender are thought to play into this pattern, reinforcing dominant ideologies and thereby reinforcing control over gender performances. Such ‘reductive’ accounts of gender delegitimize both queerness and politically subversive gender performances.

Turning to the second aspect, intersectionality entails that gender is constituted by an indefinite number of intersecting social identities. Here the issue is not that positing essences reproduces the conditions of oppression but that the experiences and life conditions of marginalized persons are so diverse across different social identities that it's impossible to find any criterion or experience shared by all and only members of any given gender. Thus any theory that purports to give a universal account of gender essence is bound to be inaccurate. Moreover, this inaccuracy inevitably produces political intersectionality (Crenshaw 1993): intersectional identities falling outside of accepted gender definitions are marginalized within political movements.

My approach to gender authenticity is designed to navigate both intersectionality and anti-substantialism. On my account, gender is to be understood (roughly) as an interpretive relationship between the individual and deeply unconscious contents, i.e. psychological contents that are never directly experienced. This inner relationship is mediated by collective cultural and epistemic resources as well as personal experiences. Each of us depends on personal, social, and historical material to interpret the meaning of gender and our own experiences of it. Gender experiences are authentic or not depending on whether they reflect the person's own interpretation of how things are at this deeply unconscious level. Authentic experiences can be fluid, non-binary, pangender, gender chill, or what have you.

There are substantial truths about gender, but these truths are found in personal interpretation. This accords with Butler's key anti-substantialist insight: theories that infringe on personal interpretive practices and create social meanings that delegitimizing authentic gender experiences and expressions are rendered problematic. Intersectionality is also accounted for. Gender is intersectional because our psychologies, our means of interpretation, our cultural and epistemic

resources, and our opportunities for personal exploration are socially situated within intersecting oppressions.

Collaborative Improvisation and Shared Intention

*Anne-Sofie Munk Autzen—Collaborative Improvisation and Shared Intention
University of Copenhagen*

Improvisation has rarely been mentioned in philosophical discussions of shared intention and joint action, despite the fact that improvisation is often highly valued as an artistic achievement in collaborative arts practices such as ensemble jamming, contact improvisation, and improvisational theatre. Improvisation is also thought important for many everyday joint actions where participants often need to adapt their responses to those of their partners or competitors on the fly. However, a number of theorists, including Paisley Livingston (2011), Garry L. Hagberg (2016), Sondra Bacharach, and Deborah Tollefsen (2010; 2011) have recently sought to account for the intentional structure of collaborative improvisation in artistic contexts. While these philosophers agree that certain cases of collaborative improvisation involve shared intentions, they disagree on which shared intention account of joint action fares best. Both Livingston and Hagberg appeal to Michael Bratman's theory of shared intention involving individual intentions whose contents interrelate in distinctive ways (Bratman 1999; 2014). In contrast, Bacharach and Tollefsen propose a perspective on collaborative improvisation based on Margaret Gilbert's plural subject theory (Gilbert 1989; 2014). I think there are problems with these two approaches to collaborative improvisation.

The first approach suggests that participants have common knowledge of rather specific intentions concerning the artistic work as a whole. For this reason, it does not adequately capture the kind of open process that is part of most forms of collaborative improvisation. The second approach requires instead that participants have an overall joint commitment to the work, and then make use of different improvisation strategies during the collaboration. Yet, this approach also relies on the idea that participants have common knowledge about what they are about to do as a plural subject. Nevertheless, there are cases of collaborative improvisation where this kind of plural subject is not in place. There is an alternative theory of shared intention that has yet to be considered as a point of departure for developing an account of collaborative improvisation: Stephen Butterfill's simple account of shared intention (2015). The aim of my talk is to argue that this theory provides an account of collaborative improvisation that successfully distinguishes between joint and parallel improvisation, and also makes sense of a broader variety of collaboratively produced improvisations than the two other approaches, because it doesn't require that roles or tasks are known by the participants in advance.

A Theory of Collective Virtue

Matthew Baddorf—*A Theory of Collective Virtue*
Walters State Community College

We say things like “Enron's greed led to catastrophe for many investors” and “the 82nd Airborne was very brave.” This language suggests that we believe that groups are capable of virtues and vices. It is hard to know what to make of this idea, however. Humans have virtues and vices due in part to our mental capacities, but attributing mental capacities to groups can sound crazy, like invoking vaguely Hegelian group spirits who work over and above their human members. Pressure to avoid such results can lead to summative views of virtues and vices: on such views, a group's virtue or vice is just a result of “summing up” the same trait in its members. (So perhaps the 82nd Airborne was brave just because most of its members were.) Unfortunately, this safely reductive view suffers from counterexamples, such as Lahroodi's example of a church committee which is closed-minded due to social pressures despite being made up of open-minded individuals (Lahroodi 2007, 287).

If we want to understand how groups can have important virtues and vices such as greed, bravery, and open-mindedness, we need a credible view that avoids the problems of Hegelianism and summativism. Here I set out such a view, a kind of non-reductive theory drawing on work in the philosophy of mind, and argue that it is worth serious consideration.

My theory consists of two claims:

Collective Virtue: (A) Groups can possess virtues and vices; (B) they can do so by possessing reasonably broad and stable dispositions that functionally and intentionally imitate individual virtues and vices.

Non-Reductiveness: Sometimes collective virtues and vices do not supervene on any set of intrinsic properties of individuals; they are thus not reducible to them.

First, I explain this view, and illustrate how it can handle a variety of cases of collective virtue and vice. In elucidating what I mean by “intentional imitation,” I go beyond previous discussions of the nature of collective intentionality, which helps me explain how groups can have virtues like ours despite their lack of genuine minds.

Then I reply to Cordell's (2016) recent sustained attack theories of collective virtue. I show that my account avoids the particular criticisms he levels against other views, and conclude by responding to his general claim that theories of collective virtue are explanatorily and practically irrelevant. Thus, I argue for my theory's superiority as a non-reductive theory of group virtue, while also defending the idea that the project of developing such a non-reductive theory is worth carrying out.

There Are Only Two Kinds of Social Kinds

*Paweł Banaś and Krzysztof Posłajko—There Are Only Two Kinds of Social Kinds
Jagiellonian University*

Our aim is to show that among entities commonly identified as „social kinds” there are exactly two separate ontological categories: *sensu stricto* social kinds and quasi-natural social kinds. The distinction between them boils down to what has been called (following Searle) “self-referentiality” of social entities.

In general, self-referentiality thesis amounts to saying that social kinds are ontologically subjective, i.e. that their existence depends on our propositional attitudes towards them and, in its stronger version, that existence of any individual member of that kind depends on someone having the relevant attitudes towards it.

Many philosophers (among them e.g. Thomasson and Khalidi) pointed out that strongly self-referential social kinds are not exhaustive of what is commonly associated with social entities: neither recession nor racism satisfy the requirement of self-referentiality, yet they both intuitively seem to be social kinds. There being a recession does not require that we all agree that indeed this is the case. That there is (or was) a recession is what social scientists (economists) may discover, even retrospectively. In many ways, as Khalidi convincingly argues, recession is like a natural kind.

Following Khalidi and Thomasson, we agree that one should admit two kinds of social kinds into her ontology: institutional (or: *sensu stricto*) and quasi-natural ones. However, our position is that these two are disjunctives, exhaustive of the more general, „*sensu largo*” category of social kinds. In this we oppose Khalidi and Thomasson, who (following Searle) suggest that there is also the third category of social kinds, composed of weakly self-referential kinds, dependent on human propositional attitudes about the kind itself but not about its members (the classic examples are war or money).

We argue that if one treats the requirement of self-referentiality seriously enough, i.e. as a constitutive requirement of existence, then there being a self-referential kind with non-self-referential members seems hardly plausible. The main argument is that for every given kind, either it is the case that its extension is determined by beliefs with a particular content, or the extension is independent of them.

The idea that there are weakly self-referential social kinds stems from the fact that at least some of the words used to name social entities are polysemous: they can refer either to an institutional social kind or a quasi-natural one. “War” seems to be a prime example here: we can be at war even if there are no warlike activities involved (but only we think this is the case); on the other hand, even if no one thinks that there is a war going on, there might be one, from a point of view of a scientist.

There is an important methodological consequence of such a dualistic approach to social kinds. What we postulate is a strong distinction between two classes of social phenomena. Institutional ones that are to remain a subject of inquiry for social ontology, whilst the quasi-natural ones should be the focus of the broadly understood social sciences.

Social Complexes, Aspects, and Many-One Identity

Donald Baxter—Social Complexes, Aspects, and Many-One Identity
University of Connecticut

Is a social complex many united people or a group entity in addition to the people? I incline to the former because I assume with Ockham that it is ontologically preferable to avoid positing additional entities if possible. For specificity, I will assume that a social complex is a plural subject in Margaret Gilbert's sense. Thus I incline to the view that a social complex is many people united by a joint commitment.

However, a metaphysical problem arises that pushes in the other direction. How it can be that, say, two individuals united can believe something that neither individual believes, while yet the two individuals united are nothing in addition to the two individuals. At first blush it seems that the two individuals united both do and do not believe the same thing, which is contradictory. To resolve the apparent contradiction, there is pressure to take there to be a group entity that has the belief in addition to the two individuals that lack it—an entity numerically distinct from each of them.

To resolve the apparent contradiction without the group entity would require distinguishing the people insofar as they are jointly committed from themselves insofar as they are individuals. That in turn would require distinguishing each person insofar as she is jointly committed from herself insofar as she is an individual. I present and motivate my theory of Aspects to argue that these differences do not require any additional numerically distinct entities. I argue that there can be qualitative complexity without quantitative complexity. That is, there is qualitative self-differing that need not—indeed cannot—be explained by saying that something is (or some things are) intimately related to or partly made up of numerically distinct things. Another way to put it is that things have numerically identical but qualitatively differing "aspects." This view may seem to violate Leibniz's Law, but I argue that it does not. Leibniz's Law only concerns individuals and perhaps pluralities but is silent about their aspects. They are not in its domain of quantification.

With the metaphysical problem resolved, I conclude that a social complex is many people united by a joint commitment and not an additional group entity.

In an epicycle, I present my theory of Many-One Identity and use it to argue that, though many united people are not an additional entity, a social complex can nonetheless count as a single entity. Thus, Gilbert is right to call them "real unities."

Collective Responsibility and Agentive Functions

Saba Bazargan-Forward—Collective Responsibility and Agentive Functions
University of California, San Diego

Proponents of collective responsibility maintain that when individuals cooperate in furtherance of a wrong, each cooperator can be responsible for more than the difference she makes. That is, causal over-determination does not vitiate responsibility in the context of cooperative activity. This view – call it ‘Strong Collective Responsibility’ – requires defense since it seems to conflict with the equally plausible pronouncement that each of us is responsible only for what is within her own causal reach. We need an account explaining how and why cooperation dramatically expands the scope of responsibility. It’s my aim to provide such an account.

I begin by addressing a fundamental issue: how must conduct be related to agency in order for an individual to be responsible for that conduct? We need a ‘bridging principle’ – a principle that provides the requisite relation between the relevant aspects of my agency on the one hand, and the conduct for which I am responsible on the other. On the account I defend, if I assign to you the agentive function of promoting ends I have selected, and you accordingly adopt that function, then your subsequent actions, to the extent that they satisfy my ends, manifest the responsibility-relevant aspects of my agency. As a result, I am responsible for what you do -- or so I argue. I call this a ‘functionalist’ bridging principle.

I assign an agentive function to you when I have intentionally made it so that my end has peremptory force for you, which is to say that my decision settles the matter for you – no further deliberation is required on your part regarding what to do. If I assign to you that agentive function, and you adopt it and act in accordance with it, then according to the functionalist bridging principle, your conduct manifests my agency, which means I am responsible for what you do.

I then analyze cooperative activity in a group of individuals in terms of multiple, symmetric, and pairwise agentive functions. Each cooperator has a local end, where that end promotes or is constitutive of a collectively caused event Ω . Each cooperator assigns to every other cooperator the agentive function of fulfilling her own local end. And each cooperator accepts such an agentive function as well. So given the functionalist bridging principle, any given cooperator will be potentially responsible for up to the sum total of what all the cooperators together do, provided that she has agreed to assist all the others in achieving their respective aims of promoting Ω . Notice that even if a given cooperator fails to contribute to Ω , she will be responsible for what her cohorts do, to the extent that they fulfill the agentive function she has assigned them (even if it is not because of her that they do so – or so I argue). Hence causal over-determination in the context of cooperative conduct does not vitiate responsibility. And this is just to say that Strong Collective Responsibility is correct.

Inside and Out: Ascriptivism About Gendered Traits

Rose Bell—Inside and Out: Ascriptivism About Gendered Traits
Syracuse University

The demarcation of gender according to intrinsic physical or psychological features has never been popular among feminist philosophers. An ascriptivist picture has greater theoretical benefit for many reasons. It explains how gender oppresses; furthermore, it avoids the pitfalls of attributing gender inequality to some intrinsic feature of the oppressed. However, some of the phenomenology reported by transgendered individuals provides a prima facie challenge to ascriptivism about gender. Some transgender individuals describe the experience of gender as a first-order, involuntary “pull”, contrary to social expectations—as if there were something intrinsic about their gender. If gender is best understood as an ascriptive social category, however, what are we to make of those who avow this intrinsic phenomenology—and can we reconcile this with the theoretical power of ascriptivism about gender?

I offer the following solution. What is ascriptively gendered, I claim, is not individuals but traits. Some such traits are physical and others are psychological. If an individual possesses the physical traits associated with a certain gender label, they are assigned this label and expected to exemplify the other associated traits. An individual’s gender is therefore not a single feature, like hair color, but rather something like a caste: it picks out some expected collection of features, behaviors, and norms.

This picture has the following theoretical benefits: 1) It explains the phenomenology reported by those transgender individuals who describe the “pull” of the gender with which they identify; what pulls on them are certain traits, gendered at odds with those they are expected to display according to their assigned gender label. 2) It achieves this while retaining the demarcation of gender as an oppressive social category. 3) It allows for a spectrum of genders, and clears the way for the ameliorative “re-branding” or “un-gendering” of traits in a way that removes their oppressive baggage.

I begin by giving a brief summary of the ascriptivist picture, and explain the theoretical appeal of treating gender as an assigned social category. I highlight the tension between assigned-category accounts and the phenomenology of many transgendered people. I suggest that gender picks out something intrinsic, but what it picks out does not correspond to biological sex. I then distinguish between describing this aspect of gender as a single trait, or as a collection of traits. I explain how the desiderata of respecting trans identities and describing gender as an assigned social category can both be achieved if we adopt the collection of traits view. I conclude by addressing the objection that gender is purely socially constructed, and does not correspond to any natural traits.

The Active Role of Technological Artifacts In Enacting Gender Norms: Collaborating With Normality

*Cristina Bernabéu—The Active Role of Technological Artifacts In Enacting Gender Norms
Autónoma University of Madrid*

One of the major philosophical concerns regarding the analysis about the gender impact of technological artifacts on the ways in which we perceive and behave has to do with determining the sort of activity they exhibit. Within the framework of contemporary philosophy of technology, articulated by the denial of the pessimistic viewpoint of technological development, the problem of technological agency has emerged as one of its most relevant issues. Robert Rosenberg (2014) has referred to it as a problem that primarily has to do with the question about whether technology reduces or help shape the choices and actions of users and, secondly, as somewhat related with the fact that technology is at once something which is constructed by “collective actors” and something that constructs such collectivity. Additionally, the active or constructive role of technological artifacts has been understood by some current philosophers of technology in the light of their potential for producing unexpected conditions and relations. Much of design ethics revolves around the moral relevance of this fact, i.e. “the moral relevance of non-human reality” (Verbeek 2009, 242).

My claim is that the role of collaborator, developed by Sally Haslanger (2012) with regard to sexual objectification processes, is appropriate to account for the kind of normative activity of technological artifacts. Such an assumption runs up against the issue concerning the agency of artifacts as it has been addressed by contemporary philosophy of technology. Particular attention is given, on the one hand, to Peter-Paul Verbeek’s proposal, which is relevant in that it challenges, through a non-instrumentalist approach, the dominant perspective inside the realm of the design ethics, which is exclusively focused on risk assessment; on the other, to the notions of Action Schemes and Second Order Responsibility, introduced by Illies & Meijers, since they open up the possibility to explain the impact of such active role on agents. I agree with the criticism of Illies & Meijers in that the responsibility of technological artifacts that results from Verbeek’s viewpoint is untenable, due to the radicalization of its moral sense. However, I claim that their proposal, which is concentrated on a “softer” influence, is too weak to account for artifacts active normative role.

From there, my main hypothesis is that, in contexts of gender hierarchy, technological artifacts have a normative impact in human action and perception in such a way that certain social norms are enacted through the exercise of human-technology interaction practices. Along with this, while certain human-technology interaction practices such as the use of smartphones, tablets and computers (just to name a few) may be very relevant for challenging traditional and hence gendered models of subjectivity, they may also contribute to reinforce “the historical structures of the masculine order in the form of unconscious schemes of perception and appreciation” (Bourdieu 2001, 5). From a feminist perspective, a critical work on how such an active role is reactivating or deactivating classical gendered roles is thus needed.

I start from two assumptions. First, I assume that contexts of gender hierarchy are to be understood in structural (both material and symbolic) terms. As a structural form of oppression, it may be defined as an “immense symbolic machine” that “tends to ratify the masculine domination” by

imposing itself as being neutral (Bourdieu 2001). Second, I assume that human-technology interaction practices must be considered as kind of social practices. Following Haslanger's recent work, "social practices are patterns of behavior that depend on learned skills and locally transmitted information, in response to resources" (Haslanger, 2017:10). Thus, from this perspective, what allows to consider human-technology interaction practices as a kind of social practices is the possibility for us to find a common pattern of behavior [by virtue of the nodes that agents occupy inside the social structure (Haslanger 2016)], in the use of certain technological artifacts.

The Metaphysics of Intersectionality

Sara Bernstein—The Metaphysics of Intersectionality

University of Notre Dame

This paper develops and articulates a metaphysics of intersectionality, the idea that multiple axes of social oppression cross-cut each other. Viewing identities in such a way has become central to understanding how the various dimensions of race, gender, sexual orientation, disability status, and class interact to form more complex forms of discrimination than those suffered by persons who fall under only one category. Though intersectionality is often described through metaphor, I suggest that rigorous theories of intersectionality can be formulated using the tools of contemporary analytic metaphysics. A central tenet of intersectionality theory, that intersectional identities are inseparable, can be framed in terms of explanatory unity. Inseparability should not be understood as modal inseparability or conceptual inseparability, I argue. Further, intersectionality is best understood as a kind of explanatory priority of the intersectional category over its constituents, comparable to metaphysical priority of the whole over its parts.

Social Acts and the Limits of Pretense

Sam Berstler—*Social Acts and the Limits of Pretense*
Yale University

What does it mean to perform a communicative act? Many philosophers have answered that we must have the correct mental states: we must have the right intentions about what we are doing, and we must have the right beliefs and desires in order to support those intentions. It helps to have a conventional system of signs, but that's merely a matter of convenience.

I'll be arguing that the traditional story has a major problem. It fails to sufficiently distinguish cases of mutually known insincerity from cases of mutually known pretense. In our day-to-day lives, we're exquisitely sensitive to this distinction.

Consider a person of devout Jewish faith, let's call her Hannah. Suppose that she is an improvisation group and, as part of a comedy sketch, says at one point, "I believe in the divinity of Jesus Christ." Now suppose that she is captured by a group of Christian terrorists, who tell her that if she doesn't say, "I believe in the divinity of Jesus Christ," they will execute her. She is willing to die rather than capitulate to saying this. We can re-run the minimal pair with a symbolic gesture instead of an assertion. Suppose that as part of a comedy sketch, Hannah is willing to participate in a mock ceremony in which she throws incense to a Druid god. But when her friends ask her to perform in the mock ceremony late at night in the woods, she isn't willing to do this anymore: this would be idol worship.

What's the difference between the cases of obvious pretense, on the one hand, and obvious insincerity, on the other? In all of the cases, everyone knows that Hannah doesn't actually believe in the divinity of Jesus Christ or the Druid idol, and she isn't inviting anyone to believe this.

Intuitively, in the sketch cases, Hannah doesn't represent it as true that she has these beliefs, but in the other cases, she would have—despite the fact that everyone knows this isn't true and that she doesn't want anyone to think it's true. I argue that we can't explain this difference by appealing to Hannah's mental state or the mental state's of the participants. Rather, we have to appeal to the surrounding context of the utterance: a sketch comedy; a ceremony (even if insincerely performed); and a conversation (even if coercive).

I call communicative acts that represent as true, regardless of the mental states of the participants, social communicative acts. In the final part of the paper, I sketch some problems for a theory of them. If social acts receive their force from context, and not from the mental states of the actors, what is the status of their "force"? In way do social acts do something, even if what they do is insincere? Finally, to what extent are social acts a matter of convention, and why would such a convention arise in the first place?

Ideology and Social Structures

Matteo Bianchin—*Ideology and Social Structures*
Università di Milano-Bicocca

Ideology and ideology critique have been recently revived as a topic in social ontology. On the one hand ideologies may play a role in (some of) the mechanisms that anchor social facts (Haslanger 2016a, Epstein 2015). On the other hand they seem to possess some puzzling features. They look both true and false, to involve both cognitive and practical knowledge, to have both a descriptive and a normative content (Haslanger 2012, 2017; Jaeggi 2009; Stahl 2016). I will defend a moderately cognitivist understanding of ideology and ideology critique, focusing on two connected questions. First, insofar as social practices are imbued with normative motives (Bicchieri 2017), a normative theory is required to ground critical claims raised on them (Haslanger 2013). Second, the normative task of ideology critique must be properly connected with its explanatory tasks in order for critique to be “immanent” and promote change (Jaeggi 2009).

I maintain that the explanatory task confronts three connected explanatory demands related to the functional, the epistemic, and the genetic dimensions of ideology (Geuss 1981, Shelby 2003). I draw on an etiological reading of functional explanations (Kincaid 1996, 2007) to explain the existence of ideology and on the epistemology of delusions and irrational beliefs to locate the distorting mechanisms at work in shaping ideological formations (Bortolotti 2010; Bortolotti 2017, Gunn, Bortolotti 2018). Finally, I draw on social structural explanations (Haslanger 2016b) to account for how social structures shape the conditions under which the relevant distorting mechanisms are triggered. On this reading ideology is located at the level of the mechanisms that anchor social facts. It rationalizes domination by manipulating the rationale for accepting the social practice they are designed to support and may eventually enter social structures by virtue of looping effects, thus reinforcing the conditions under which distorting mechanisms are triggered. At this stage ideology may operate as an embodied set of practices that need not involve belief and can be part of the feed-back mechanism that explains the persistence of ideology (Bourdieu 1977). What licenses the idea of ideology critique, however, is that the social function and the (purported) normative content of ideology are connected, as ideologies perform the former by virtue of the latter, so that their normative shortfalls are bound to result in legitimation crises that unsettle social cooperation (Jaeggi 2016; Habermas 1988).

I argue that the explanation of ideology hence involves the idea that ideology fails to provide principles to regulate social cooperation which would be accepted under conditions of non-domination, thus failing to anchor a cooperative scheme that is stable over time. In this sense, by explaining ideology as parasitic on domination, critique both debunks the purported reasons it provides for accepting existing social practices and points to the conditions under which cooperation stabilize as those of a social practice whose principles would be accepted without ideological coercion. Thus, ideology critique seems to entail a normative theory that comes close to a theory of justice designed to articulate such principles as part of a theory of social cooperation.

The Dynamics of Norm Erosion

Cristina Bicchieri—The Dynamics of Norm Erosion

University of Pennsylvania

We study how exposure to peer behavior affects the dynamics of norms. In a novel multi-period and non-strategic setup, individuals can actively comply with or violate a pro-social norm of giving. We study norm erosion by varying the observability of both peer behavior and social proximity among peers across treatments. Overall, we find that exposure to peers drives the erosion of norms by facilitating the spread of norm violations. In the presence of social proximity to one's peers, however, individuals are influenced by observing both examples of norm violations and norm compliance. We also find substantial heterogeneity with respect to how susceptible individuals are to behavioral change. These insights are particularly important from a policy perspective because they can inform and improve the effectiveness of norm-based interventions at the individual and collective level and help to advance our understanding regarding the role of social proximity (identity) in the dynamics of norms and behavioral change.

The Importance of Shared Obligations

*Gunnar Björnsson—The Importance of Shared Obligations
Stockholm University & University of Gothenburg*

Elsewhere, I have defended the idea that a group of agents can have moral obligations even if it does not itself constitute a moral agent. On the proposed account, obligations of the relevant sort are ultimately based on demands that individual moral agents care about—are disposed to promote—morally important matters to certain degrees. More specifically, our obligations are behaviors that would be ensured, in normal ways, by episodes of required caring, and the absence of which would be morally undesirable. In the case of individual obligations, the relevant behavior is ensured by one or more episodes of required caring, which guarantees that the agent performs or omits some action, or brings about or allows for a certain outcome: helps someone in need; abstains from lying; preserves a relationship, say, through a series of individual actions or omissions. In the case of what we might call "shared" obligations, episodes of required caring in more than one individual similarly ensure some behavior or outcome: that two individuals preserve their friendship over time, say, or that a group of teenagers let the neighbors sleep throughout the night.

Given that the relevant sort of obligation is ultimately based on requirements that individual moral agents care about various values, it is natural to ask why shared obligations matter. Arguably, ordinary moral thinking does in fact attribute obligations to non-agential groups. But what important theoretical or practical role do shared obligations play? Why not just focus on the demands on individual caring that ground these obligations?

In this talk, I outline two such roles.

First, obligations of the relevant sort correspond to a general interest that we have, as social moral agents, in how important outcomes are related to what agents do and do not care about. Following traditions tracing back to Hume and Strawson, I take our pervasive practices of blaming to bear witness to this interest, as blame is paradigmatically prompted by cases where substandard caring has morally undesirable upshots. As understood here, obligations directly relate to the very same interest. Moreover, since morally important outcomes often depend on the degrees of caring of several individuals, this interest leads us to consider shared as well as individual obligations.

Second, shared obligations play an important role in a pervasive sort of practical reasoning: "issue-oriented" reasoning. Such reasoning starts with a practical issue—some good or bad possibility might be promoted or prevented. It then proceeds to identify possible means to such promotion or prevention, means that may involve one or more agents, and to evaluate these means. Such evaluations will often be sensitive to whether the means would be taken by appropriately caring agents, i.e. to whether taking them is an obligation in the relevant sense. Moreover, appropriate caring will itself often ensure that individual agents engage in issue-related reasoning that considers group behaviors as relevant means and considers whether they are shared obligations. On the proposed account of obligations, this means that we often have individual obligations to consider shared obligations.

Are Mass Shooters a Social Kind?

*Kurt Blankschaen—Are Mass Shooters A Social Kind?
Boston University*

On April 20, 1999, Eric Harris and Dylan Klebold went on a shooting spree at their high school in Columbine, Colorado. Although Harris and Klebold weren't the first to open fire in a school, national media reports dubbed the event a "school shooting." The term stuck and grimly expanded to "mass shooting" to account for similar shootings in other public places. Politicians, psychologists, law enforcement agencies, and school safety policies anchor mass shooters as a particular kind of threat, constituting a particular social category. We should disabuse ourselves of this framework for two reasons: first, it is descriptively vague and not an effective way to talk about who counts as a mass shooter or what counts as a mass shooting; second, reactionary policies and laws generate negative moral externalities.

I start by identifying two ameliorative reasons why someone might want to treat "mass shooter" as a social kind and then show why each of these fails. First, if we can identify a cluster of character traits of people who are at risk for becoming a mass shooter, then we can use a predictive profile for intervention strategies. One commonly circulated cluster theory is: young, white, men, afflicted by mental illness, and who have access to firearms. Since every trait does not have to apply to every case, this kind of a cluster theory accounts for paradigm cases of mass shooters, while still including cases where the shooter does not have all of the traits. This cluster, however, overcommits us because if someone has every trait, then they count as a member of the category. Too large of a demographic ends up counting: any young, white male with a mental illness and access to firearms would be at risk for being a mass shooter. The predictive profile loses its preventative value.

Second, if we can distinguish mass shootings from other kinds of crimes, then we know if our policies or laws are effective. Body count is a bleak, but objective, standard. If, however, we set the body count too high, say, three, then we exclude shootings with many injuries, but few deaths. If we set the body count too low, say, two, then we artificially inflate the number of mass shootings by counting crimes, like a robbery gone bad, that are not mass shootings. These arguments should give us pause about how predictively effective using "mass shooter" as a social kind is. Many thick social kinds, however, are descriptively vague, but normatively important. "Mass shooter" is different because it perpetuates a moral panic. Framing mass shooters as social monsters "out there" requires us to "do something." Reactionary policies and laws are often poorly conceived, but retain an aura of protection. Heightened security measures in schools or zero-tolerance policies contribute to higher rates of expulsion and lower graduation rates, especially for students of color. News reports that emphasize mental illness falsely associate mental illness with violent or lethal outbursts, making it harder for people to seek help or even publicly talk about their mental illness.

Responsibility Gaps and Attributionism

Olle Blomberg—Responsibility Gaps and Attributionism

Lund University

Following an explosion on Deepwater Horizon in 2011, massive amounts of oil leaked into the Gulf of Mexico, with devastating ecological and economic consequences. Investigations primarily blamed BP, the operator of the rig, for the disaster. The causes included mistakes and failures by many individuals and teams. Cases such as these arguably present us with so-called "responsibility gaps": After having held all the individuals involved responsible, there seems to still be some responsibility unallocated (so to speak). Philip Pettit have argued that such gaps can be closed in some cases where the larger organisation (BP, in this case) is an appropriate target of responsibility attributions (Copp 2006; Pettit 2007, 2017). In this talk, I argue though, that Pettit's intentional control-based account of moral responsibility is too demanding to make sense of many possible genuine responsibility gaps. In particular, Pettit's conditions that an agent must have "Full knowledge of guilt" and "Full consent of the will" (2017, p. 28) in order to be fit to be held responsible won't be satisfied in some cases where we nevertheless appear to have a responsibility gap. Pettit argues that cases where these conditions aren't satisfied will be uncommon (2007, pp. 186-87; 2017, pp. 31-32), but the question remains how such cases, even if uncommon, should be treated. I argue that a less demanding "attributionist" view of moral responsibility (akin to that favoured by, for example, Tim Scanlon, Angela Smith and Matthew Talbert) is better positioned to make sense of such responsibility gaps that involve organisations that are agents but lack full knowledge of guilt as well as full consent of the will for the wrongdoing in question. This is not very surprising since this type of attributionist view is best positioned to make sense of analogous individual responsibility gaps where an agent's actions or omissions cannot be traced back to some knowing wrongdoing in the past.

Solving Lewis and Hart's Problems with Normativity, Evolution, and Disagreement

Sarah Braasch—Solving Lewis And Hart's Problems with Normativity, Evolution, And Disagreement
Yale University

David Lewis and HLA Hart have very similar problems with their accounts of social conventions in *Convention and the Concept of Law*, respectively. They both want to capture an obligatory character of social conventions. They both think that the only way to capture this obligatory character is if there is a single rational thing to do in a situation that gives rise to a social convention, which is to conform as everyone else in one's population conforms. They both think that if there is a single rational thing to do, then you have an obligation to do that one rational thing. But, these Lewisian social conventions only arise in situations wherein everyone is particularly motivated to coordinate their behavior and achieve unanimity of conformity. The problem with Lewisian social conventions is that they must pop into and out of existence instantaneously, and they can neither evolve nor devolve over time. Nor do they allow for disagreement or pluralism. This is the price that one must pay to capture the obligatory character of social conventions. But, we know that social institutions evolve and devolve over time. People disagree about what the law is and should be. And, despite this fluidity and pluralism, we still speak as if there were a real obligation to follow the law.

The first step to solving this problem is to adopt Margaret Gilbert's insight that social conventions are social group constituting. The fact that a population has a social convention constitutes that population as a social group. The second step is to recognize the role that practical authority plays. This is the case, even for Lewisian social conventions, because of the risk dominance of the status quo position. No one is going to conform to an alternate Lewisian social convention, unless she knows that her entire community is going to do likewise. The worst possible outcome is for anyone to fail to coordinate. Practical authority solves this problem, because the authority makes it known how everyone in the social group will behave. The third step is to recognize that all social conventions are step public social goods. This is the case because the authority is a freeriding defector. Now, we no longer require unanimity to generate normativity. Once the step public social good of the social group itself has been generated, the social group rests upon an equilibrium point. No one wants to defect, because the social group would collapse back to the status quo. And, the authority has no incentive to conform, because she fares far better by continuing to defect. Thus, there is a single rational thing to do in this situation, which is to continue to conform to the social convention. Now that we no longer need unanimity to achieve normativity, we can have sub social groups within a larger social group with their own sub social conventions. Thus, we can have evolution and disagreement. And, we have solved Lewis and Hart's problems.

Institutional Functioning, Institutional Agency

Michael Bratman—*Institutional Functioning, Institutional Agency*
Stanford University

Bratman (2014) develops a planning model of small-scale shared intentional activity, shared intention, and shared policy. Can this model be scaled up in theorizing about larger social institutions?

Distinguish two questions: (1) Can this planning theory help us model the functioning of an important kind of institution? (2) Will such institutional functioning normally constitute institutional agency? (This question bears on issues about accountability not discussed here.)

Begin with ideas broadly from Hart (1961): An important kind of social institution involves “social rules”, primary and secondary; these social rules have an “internal aspect”; they include secondary rules of procedure that specify how to settle basic practical issues; social rules may be accepted only within a special sub-group (e.g., in Hart’s theory, the legal officials) but this acceptance may support wider-ranging conformity; in settling relevant practical issues these secondary rules can induce institutional commitments/intentions.

What is a social rule? Draw on the planning theory of shared policies. A social rule in population P involves a background shared policy of conformity, a policy that is shared at least within a core sub-population and the sharing of which can induce reasons (e.g. via a system of incentives, or via norms of authority or fair play) that support a broader conformity to the rule within P. The policy-acceptances, on the part of individuals, that help constitute the shared policy constitute the internal aspect of the social rule. A social rule, while grounded in a relevant shared policy, is not identical with that shared policy, since the range of the social rule can include those who do not participate in the underlying shared policy but conform rather because of further reasons induced by the shared policy. Further, institutional intentions that can be the issue of a procedural social rule need not be shared intentions (in the sense of the planning theory) on the part of the overall population. And the conjecture is that we can usefully model the functioning of an important kind of social institution by way of such a merger between Hart and the planning theory. (Compare/contrast with Shapiro (2011).)

What about institutional agency? The cited institutional functioning will not in general be a shared intentional activity of the overall population. Nor need it involve the overall, holistic coherence of attitudes that Davidson and others highlight in individual intentional agency. Nevertheless, we can turn to a generic model of intentional agency. Roughly: intentional agency involves activity organized and guided by inter-related intentions, where this proceeds by way of (a) sensitivity to norms of plan rationality, (b) inducing synchronic and diachronic broadly-Lockean ties within the agent, and (c) thereby helping to constitute a standpoint that specifies a standard of success and of inclusion/non-inclusion of sub-activities. The cited institutional functioning can be such intentional agency of that institution even though it is neither the shared intentional agency of the population nor the agency of a single unified subject. It can be institutional agency without being individual or shared intentional agency writ large.

Social Contradictions

Thomas Brouwer—*Social Contradictions*
University of Leeds

Any theory of social facts has to somehow reconcile their human malleability – the way that such facts are constitutively, not just causally, shaped by our attitudes and behaviour – with their mind-independence and objectivity. The more plausible theories do this by giving a two-factor analysis, which has social facts arising from the interaction of a human, often intentional element with a typically mind-independent ‘worldly’ element. In Searle’s (1995, 2009) model, social facts arise through the imposition of collectively accepted/recognised ‘constitutive rules’ on physical reality. In Einheuser’s (2006) model, conventional reality (subsuming social reality) results from the refraction of a worldly ‘substrate’ through a conventionally-fixed ‘carving’. In Epstein’s (2015) model, the most developed of its type, social facts are grounded in non-social facts by means of contingent grounding principles, which are ‘anchored’ in place by human behaviour and intentionality, broadly construed. I refer to these approaches as two-dimensional social ontologies.

It seems to have gone without discussion in the literature that such two-dimensional models allow, in principle, for mechanisms whereby types of logical anomalies which we typically associate with human thought and talk, such as inconsistency, are translated into mind-independent reality (though cf. Priest 1987, ch. 13). For instance, a community may (perhaps through lapses in collective rationality) fix in place grounding principles/constitutive rules/carvings which, given certain inputs, yield contradictory or indeterminate facts. Of course, such possibilities may be only *prima facie* ones, for in response one might refine one’s theory of social reality by means of monster-barring amendments. Things get philosophically interesting when we ask whether such amendments could be well-motivated, both in the sense of not being *ad hoc* and in the sense of being necessary at all.

In this talk I discuss the theoretical possibility of incoherent social realities, focusing in particular on the provocative case of inconsistencies (‘social dialetheias’). After setting out the *prima facie* case for social dialetheias, I consider how one might dispute that case, and assess the balance of considerations that emerges. I then discuss whether it would be viable to embrace the possibility of social inconsistency as simply a surprising feature of social reality, by reviewing some reasons why dialetheias have been considered philosophically unacceptable. In addition to well-worn logical and metaphysical worries about dialetheias, social dialetheias may give rise to specific worries given that social facts often have a normative upshot: would social inconsistencies yield irresolvable normative dilemmata, violating *ought-implies-can*?

Finally, I discuss how a community might respond to the discovery of a social dialetheia by revising their grounding principles/constitutive rules/carvings, and how, on a two-dimensional model, such revisions could amount to retroactive changes (‘retcons’) in social reality.

Telic Power

Åsa Burman—*Telic Power*
Stockholm University

Institutional fact and deontic power are two key concepts in social ontology. I argue that deontic power, however, is too narrow to capture a central dimension of the social world, exemplified by gender and class norms. In her case study on respectability, sociologist Beverley Skeggs refers to the norm of being a good housewife as having an impeccably clean home, respectable clothes, and refined language. One of the interviewed women shares her thoughts after a Health Visitor has inspected her home:

"You know, they're weighing you up and they ask you all these indirect questions as if you're too thick to know what they're getting at and you know all the time they're thinking 'she's poor, she's no good, she can't bring her kids up properly' and no matter what you do they've got your number. To them you're never fit, never up to their standards." (Skeggs, 1997, p. 3)

The woman mentions a standard that she fails to live up to in the eyes of the Health Visitor. This example illustrates one sense of a "norm" as an existing ideal, or standard, which objects and individuals can be measured against. These ideals point to another aspect of our various social roles; an aspect which cannot be analysed by using only the concepts of deontic normativity and deontic power. To capture this aspect, I introduce another category of power – telic power – which has hitherto been overlooked:

TELIC POWER: An agent A has telic power in domain D if and only if there exists an ideal such that A can be measured against and the perceived distance of A from the ideal, by other agents, affects A's ability to effect certain outcomes in that domain.

Using examples of gender and class norms, I show that it is theoretically useful to make a distinction between positive and negative telic power, understood as a variation of the general formula of telic power:

POSITIVE TELIC POWER: An agent A has positive telic power in domain D if and only if A is perceived as living up to the ideal, i.e. she is viewed as a good exemplar of the relevant kind, by other agents, and this positively affects, or enhances, A's ability to effect certain outcomes in that domain.

NEGATIVE TELIC POWER: An agent A has negative telic power in domain D if and only if A is perceived as not living up to the ideal, i.e. she is viewed as substandard or as a bad exemplar of the relevant kind, by other agents, and this negatively affects, or restricts, A's ability to effect certain outcomes in that domain.

I end by giving two reasons for introducing this new concept: First, it captures a distinct central dimension of the social world, previously neglected due to the one-sided use of examples and the subsequent emphasis on deontic power. Second, it is theoretically useful since one can explain how telic power can both conflict with and reinforce our deontic powers.

The Role of Motivation in Structural Explanations of Oppression

*César Cabezas—The Role of Motivation in Structural Explanations of Oppression
Columbia University*

In recent years, philosophers have sought to explain the durability of systems of oppression in terms of structural explanations of why members of oppressed groups contribute to their own oppression (Cudd, Haslanger, Ayala). One prominent example is the tendency among women to quit their jobs after having a baby in order to focus on childcare. This tendency has the effect of reinforcing power asymmetries in their relationships and to exacerbate gender inequality in the labor market insofar as it encourages employers not to trust women to “stay in their careers or that if they do, they will devote the kind of time and energy to them that men will” (Cudd 2006). Whereas an individualist explanation of this phenomenon may focus on women’s “natural” nurturing tendencies and/or their free choices, a structural explanation shows that women’s choices are not fully free insofar as the social structures in which they are enmeshed constrain their possibility space for agency (viz. absence of affordable childcare, sexist norms about what it means to be a “good” mother, a gender wage gap that makes it “rational” for the couple to choose that the man will keep his job and the woman will quit hers, etc.). Thus, structural explanations shed light on how members of oppressed groups may contribute to their own oppression because of a structural constraint on their agency.

While sympathetic to this project, this paper seeks to expand the current understanding of how structural explanations feature in explanations of oppression. To do so, I suggest that we theorize structural explanations not only in terms of structural constraints on agency, but also in terms of the motivational force that social structures have on human agents. Moreover, I propose that structural explanations of oppression shift the focus away from the victims of oppression and towards those who benefit from it. As Tilly (1998) shows, a crucial feature of a structural explanation of durable oppression based on categories such as gender, race, and ethnicity has to do with practices of social closure, exclusion and control among privileged groups. In order to explain the resilience of such practices, I develop a motivational account of structural explanation that highlights the psychological mechanisms whereby social structures motivate members of privileged groups to participate in practices that reproduce an oppressive social order. Drawing on the psychological literature on group position theory, I argue that, as a result of their privileged position in unjust social structures, members of dominant groups develop a sense of proprietary claim over certain rights, statuses and resources, which motivates them to engage in practices of social closure, exclusion and control in order to maintain their privileged position relative to subordinate groups—especially when there is a perception of threat from a subordinate group encroaching on their perceived prerogatives. Importantly, this structural explanation is not only rooted in rational choice theory—as seems to be the case with Haslanger and Cudd—but also driven by affective elements (i.e. feelings of threat).

Hegelian Collective Rationality, Recognition, and International Business

*John Cantwell and Tobey Scharding—Hegelian Collective Rationality, Recognition, and International Business
Rutgers University*

We set forth an interpretation of Hegel's notion of collective rationality and show that it offers a model for ethical conduct in the system of production. We focus on production in international business (IB) in particular. Our interpretation draws on Wood (1990) and Williams (1992), arguing that rationality (i.e., as it underwrites the system of production) is fundamentally a social phenomenon. As such, individual rationality is always incomplete. Whereas the association between individualism and rationality has been entrenched in the English-speaking world since the Enlightenment, the complexity of IB calls for a more capacious understanding of rationality. In particular, the contemporary, networked context of IB corresponds well to the Hegelian concept of recognition: the "I" exists stably in the social world only when, and to the extent that, other people recognize it.

We draw two upshots from this interpretation. First, apparently individual decisions have an inextricably social component. Second, and more directly relevant to our project, decisions that have traditionally been attributed to individual rationality—e.g., those associated with multinational enterprises (MNEs)—depend on social recognition. This dependence begins, but does not end, in the decisions' ethical legitimacy. Drawing upon recent analyses of the importance of networks in social evolution and productivity (Henrich, 2016; Christakis and Fowler, 2009; Hutchins, 1995), we argue that as increasingly dense social networks make people more aware of ethical problems in IB, people's tolerance for unethical business activities (and the productivity that depends on them) diminishes. In outcries over Western MNEs' (a) treatment of foreign sweatshop workers, (b) adherence to the lower pollution standards of developing countries, (c) decisions to adhere to foreign countries' (free-speech controverting) demands for censorship, and so on, we show that the controversies cannot be solved without a conception of collective rationality that attends to the many conflicting voices at play.

The value of our analysis is a deeper understanding of the Hegelian basis of collective rationality (along with related notions of collective intentionality, group cognition, and collective responsibility). The upshot of our analysis, in turn, is an ethical paradigm for international business. In particular, our ethical paradigm has value in circumstances in which MNEs encounter contradictions in internal decision-making. An MNE must present a variety of faces to the world, some of which are evidently contradictory if one contrasts them directly; it must also reconcile a diversity of potentially conflicting interests or stakeholders. The decisions of Western MNEs discussed in the foregoing paragraph exemplify these conflicts. While controversies associated with many such decisions are ongoing, we argue that an MNE whose decision making is informed by the Hegelian notions of collective rationality and recognition can reconcile contradictions while (selectively) preserving opposites, thus supporting organic change and progressive development in the nature of the MNE—and its wider positioning in the social world.

Disability Studies, Conceptual Engineering, and Metalinguistic Negotiation

Elizabeth Cantalamessa—Disability Studies, Conceptual Engineering, and Metalinguistic Negotiation

University of Miami

Disability studies is an interdisciplinary field aimed at understanding the complex nature of disability and its related concepts, norms, and practices in order to improve the lives of people with disabilities. A crucial step in disability activism and scholarship was the advent of the “social model of disability” which holds that disability is a wholly social phenomenon imposed on physical impairment. That is, it is strictly social causes that are involved with the negative experiences associated with having an impairment. Therefore, it is the duty and obligation of society to fix the harms associated with being disabled. However, according to philosophers such as Elizabeth Barnes (2009), claims made by theorists working in disability studies appear to contradict both academic and common sense beliefs about the nature of disability. I will argue that such claims, regardless of whether they are true, are pragmatically doing something reasonable, significant, and continuous with core debates within mainstream analytical philosophy. Specifically, disability studies has been engaged in and successful at what philosophers like Burgess and Plunkett (2013a, 2013b), Thomasson (2017), Cappelen (forthcoming), and Haslanger (1999, 2000, 2004, 2012) call “conceptual ethics” and “conceptual engineering.” Broadly speaking, conceptual ethics involves normative or evaluative conceptual analysis and critique. Conceptual engineering involves the unintentional and intentional change in meaning for certain terms or concepts for particular purposes (e.g. to combat oppression). What concept(s) applies to a certain phenomenon has implications for how we understand and interact with the world. According to Plunkett and Sundell (2013), and Thomasson (2017), many disagreements over conceptual frameworks actually take place more or less implicitly. They diagnose such debates as “metalinguistic negotiation.” Cappelen (forthcoming) has argued that conceptual engineering, as presented by Haslanger, Plunkett, and Thomasson, is at best incredibly messy, complex, and uncontrollable or, at worst, impossible due to the metasemantic constraints of externalism. I will show how debates and issues from disability studies provide illuminating data for theorists interested in conceptual engineering that can dissipate some of the concerns raised by Cappelen.

This paper will proceed as follows: First I discuss the emerging fields of “conceptual ethics” and “conceptual engineering.” I then discuss the role of “metalinguistic negotiation” in conceptual ethics and engineering. In the second section I describe and explain the field of disability studies, focusing on how disability activism and theory has enabled change in the social and legal meaning of ‘disability’. In section three I argue that disability studies fits the conceptual ethics/engineering framework, and that pervasive debates from within disability studies over the proper model of disability exhibit the markers of metalinguistic negotiation. I conclude by suggesting that disability studies provides an illuminating example of successful work in conceptual engineering that is continuous with mainstream analyses in analytic philosophy.

Group Ontology and the Right Group Epistemology

*Will Zhang Chen—Group Ontology and the Right Group Epistemology
National University of Singapore*

This paper warns group epistemologists against borrowing from individualistic epistemology without due attention to group ontology. In particular, I argue that familiar diachronic norms are generally inapplicable to the group case because of important ontological differences between individuals and groups. This presents a problem for at least two recent major proposals for reliability-based norms by List & Pettit (2011) and Goldman (2014). The applicability of diachronic norms in individual epistemology is contingent on two key assumptions about individual agents. These assumptions are about numerical persistence and doxastic continuity during the paradigmatic agent's lifetime. If these assumptions do not hold, a diachronic norm like reliability will fail to live up to its precepts against epistemic luck. Once we pay attention to how group agents and their beliefs are, we should see that the key assumptions do not play out the same way. Social ontologists provide the insight that groups are something over-and-above mere sets of individuals. From there, the right group ontology will reveal other essential properties of group agency that explain why groups are persistent without doxastic continuity. Without attention to group ontology, List & Pettit and Goldman's proposals fail to grapple with these group features that render their reliability-based norms inappropriate. In fact, intuitions in favour of group diachronic norms more generally can be systematically explained away. Variants of diachronic norms across the externalism-internalism divide face a similar cluster of problems. Instead, we have reasons to think that synchronic norms turn out exactly suitable in group epistemology, not merely second-best as a retreat from diachronic norms. Synchronic group norms can combine with diachronic individual norms without the interference of diachronic group norms. Normative theory would then be more elegant. My discussion does not rule out possibilities where synchronic group norms are not quite enough for all intended our normative work. If diachronic group norms must somehow re-enter the fold, they are unlikely to be familiar extensions from individualistic epistemology.

The Social Ontology of Panajotis Kondylis: Metaphysics of the Social Being

Panagiotis Christias—The Social Ontology of Panajotis Kondylis: Metaphysics of The Social Being

Univdersity of Cyprus

The last work of Panajotis Kondylis (1943-1998) is the first of the three volumes on a theory of Social Ontology: *Das Politische und der Mensch: Grundzüge der Sozialontologie. Bd.1. Soziale Beziehung, Verstehen, Rationalität*. Berlin, Akademie Verlag, 1999. This ontology is based on the understanding of the polemical and strategical nature of all social interaction. The central idea of this scheme is the existence of a continuous "spectrum of social relations" (Spektrum sozialer Beziehungen) from extreme friendship (self-sacrifice) to extreme hatred (murder, killing). Social interaction is a mind reading process of anticipation of the action of the other, in a constant repositioning within the spectrum: living with the other is at the same time living against him. This process includes all human activities, from politics to war, and from associative to entrepreneurial life. Social life begins when two unite against a third, notices Kondylis. I would like to present, trace the origins and search the limits of this monumental essay to describe human action.

Persistent Problems In (Some) Social Scientific Conceptions of Property and Law

*Daniel Cole—Persistent Problems In (Some) Social Scientific Conceptions of Property And Law
Indiana University Maurer School of Law & School of Public and Environmental Affairs*

This conference paper draws upon two prior papers which, between them, examine three discrete, persistent problems in social-scientific applications of legal doctrines: (1) some economists treat the fundamental legal attributes of property quite loosely, supposing the existence of property “rights” where they plainly do not exist; (2) some social scientists have developed conceptions of property law based on roman law typologies, but their understandings of the roman law types often differ in very important ways from how roman lawyers understood those categories, resulting in confusion; and (3) some scholars from other social-science disciplines pay too little attention to formal legal rules, including rules of property law, on the fallacious presumption that formal rules are generally irrelevant to actual social relations. After assessing each of those three problems, the paper offers simple and straightforward (but not novel) means for avoiding them, and offers a simple tripartite typology of relations between formal and informal legal systems and rights: (1) formal legal rules sometimes (probably not an insignificant number of cases) are the working rules; (2) in many cases, the working rules are the formal legal rules as amended by social norms; and (3) in some (probably relatively few) cases, the formal legal rules are completely unrelated to the working rules. The paper offers several examples of each category, and explains how understanding relations between legal rules and social norms, along with legal understandings of concepts such as possession, can help social scientists avoid both unwarranted assignments of “property rights,” and unwarranted discounting (or outright dismissal) of legal rules, including property rights.

Blameworthiness for Non-Agent Groups

Stephanie Collins—Blameworthiness for Non-Agent Groups

Australian Catholic University

This paper argues that groups such as 'humanity', 'carbon emitters,' 'misogynists,' and 'white people' can be blameworthy for the ills they cause—even though such 'non-agent groups' cannot bear irreducible obligations not to cause those ills. This is because blameworthiness functions as a reflection of the esteem or disesteem with which others should hold an entity that has produced some outcome, whereas obligation function as an input into the reasoning of the entity that bears the obligation. Whether blame is fitting in a (type of or token) case is determined by reflective equilibrium between (1) considered convictions about similar cases; (2) the consequences of blaming in that type of case; and (3) the fairness of blaming in that type of case. I use the recent #metoo campaign to demonstrate that these three conditions sometimes make a non-agent group apt for blame. I propose three conditions that are jointly sufficient for a non-agent group's being blameworthy: (1) the members of a non-agent group each hold a particular attitude (perhaps implicitly); (2) their actions based on that attitude combine to create or uphold webs of norms and expectations that reinforce the attitude and the actions performed on its basis; (3) a harm results from the combined actions, where we cannot say precisely which members did precisely which wrongful actions to produce precisely which aspects of the harm.

On the Distinction Between Natural and Social Kinds

Rachel K. Cooper—On the Distinction Between Natural and Social Kinds
University of California, Irvine

In this paper I attempt to carve out an original account of the metaphysics of social categories such as gender and race. I begin by giving a brief summary of what I take to be some of the most important insights within the debate thus far, paying particular attention to Social Construction views. I then give reasons for thinking that social constitution views actually cloud important aspects of the phenomenon in question and that a byproduct of this approach is an unnecessary alienation of philosophers of many stripes. I offer a view that is thinly constructionist, uses ordinary rather than perceptual individuation, and instead captures the social forces at play with the concept of a mimicked disposition. I argue that it has greater epistemic benefits and therefore greater potential political benefits than social constitution views in capturing ordinary intuitions while still opening them to critique, ruling out generalizations about group members in a straightforward way, and avoiding an unnecessary bifurcation of the ontology into two different types of kinds.

Institutions, Functions, and Ethics

Sean Cordell—Institutions, Functions, and Ethics
Open University, UK

There is a cluster of problems in applied ethics and social philosophy, solutions to which require explanations and evaluations of social institutions themselves rather than just the conduct individuals. For example, how binding on an individual is this or that putative obligation of some institutionally defined social role, and why? Which particular obligations should or should not be attached to which institutional roles, and why? When and why has some institution somehow 'gone wrong' in its circumscribing particular duties of its offices? To adequately address these kinds of questions is to ask, at the most basic level, why and in what forms particular social institutions (should or should not) exist. I argue that this strongly invites formulation of some functional explanation in terms of: 'social institution x is there to do y'; and a functional evaluation in terms of 'x is a good instance of its kind when it performs y'.

Social institutions seem especially apt for such functional talk insofar as they appear to be social artefacts. As for a 'what it is there for' type explanation, if we suppose that institutions are or were 'designed' for some purpose, then we might suppose that their proper functions be as intelligible as those of tools such as hammers or pens. A 'good-of-its-kind' type functional evaluation then appears to follow from this ('a good school teaches its students well'). In addition, institutions are or should be 'for' the good. That is, insofar as they are or should be pro-social entities, institutions should instantiate some good in the world. (Consider that 'failing school' does double work: failing by its own standards qua school, and failing those who should benefit from it.) Yet for all its aptitude and potential utility, a satisfactory functional account of institutions proves elusive. We can for example find a hammer being used to do something other than hammering, or serving some non-hammering purpose, and say it is failing to fulfil its proper function. But this is nothing like as straightforward in the case of the school, the university, the family, the church, or the Senate for example. There are several reasons for this. One is that the 'design function' of an institution is in many cases hard to identify or to separate from its current or recent historical 'use function'. Another is that even where an original purpose or intention is clearly identifiable as that for which some institution was created, it may no longer play any part in that institution's current putative proper function, and for good reasons.

I explore these, and other associated, problems by considering some candidates for functional analyses and adapting them to the constitution of social institutions. Drawing on their pros and cons, and eventually revisiting Aristotle, I suggest in conclusion that: An institution's function (characteristic activity) is that it serves specific social good(s) in a way which makes it that kind of institution and not some other. I end by discussing some qualifications to, and limitations of, this proposal.

Amelioration and the Ethics of Gender Ascription: Why Changing Our Concepts Isn't Enough

*Kevin Craven—Amelioration and the Ethics of Gender Ascription
University of Michigan*

Social constructionists often find themselves at odds with folk intuitions about the meanings of socially significant terms. Whereas common sense arguably has it that one's gender is determined by one's anatomy, constructionists often claim that gender is a fundamentally social phenomenon. Such clashes with common sense may be integral to constructionist theories' radical potential, but they also raise a worry about failure to meet people where they are. Insofar as folk intuitions determine the meanings of the relevant terms, the counterintuitiveness of constructionist analyses may suggest that constructionists are 'changing the subject,' referring to phenomena other than those under discussion in most discourse on gender. Sally Haslanger famously responded to this worry by arguing that social constructionist analyses could be understood as aiming to capture not (just) the concepts we currently possess, but rather the concepts we should possess. Haslanger dubbed this sort of normative inquiry into concepts amelioration. The idea has since become widespread in feminist philosophy.

I argue that amelioration, as it is usually understood, fails to adequately respond to the spirit of the subject-changing worry. By leveling its critique at the metalinguistic or metaconceptual level, it fails to engage with first-order, ground-level disputes about the boundaries of gender groups as they actually are here and now. This is brought to the fore when we consider contemporary popular debates over bio-essentialism and the legitimacy of trans identities. Opponents of bio-essentialism claim not just that trans women ought to be labeled or viewed or treated as women, but that they are women. Amelioration as conceptual revision cannot do justice to this claim.

I sketch a solution to this problem predicated on a fundamentally pragmatic account of the meanings of gender terms. I argue that a satisfactory account of gender terms' meanings should aim to capture social functions rather than extensions. Such an account could allow us see the popular debate over bio-essentialism as not a clash of meanings, but a first-order dispute occurring within gendered practice. On this picture, the worry about 'changing the subject' largely disappears. The development of a counterintuitive account of gender can be viewed as a political act deploying the pragmatic meanings of gender terms as they currently are, rather than as a proposal calling for those meanings to be revised.

Plural Subjects and the Question of Class Consciousness

Ben Curtis—Plural Subjects and the Question of Class Consciousness
University of Memphis

In her book *A Theory of Political Obligation*, Margaret Gilbert develops a theory of plural subjects. In it, she argues that an economic class does not constitute a plural subject because the individuals that constitute it are not committed to act in the right kind of way, that is, jointly commit as a body. This may be true in our current historical moment, however from a Marxian perspective it does not fully capture what an economic class is. Today, the members of the working class, particularly in the United States, may not have jointly committed to anything as a body (yet), but there is no reason in principle why it could not, in the future, do so. Further, while Gilbert's account accurately captures the subjective aspect of what class consciousness might look like, it does not address the objective aspect elaborated by Marx and his subsequent commentators. Objectively speaking, a class is constituted not by its joint commitment to act, but rather the material conditions and organization of society. For Marx, members of the working class are bound together in a class even if they are not aware of that class membership. Class antagonisms, not joint commitments structure this aspect of (potential) class consciousness. To develop a full or robust account of class consciousness, one should take both subjective and objective aspects into account. The relationship between these two aspects is vital to understanding the potential transition from an 'unconscious' economic class to an economic class fully conscious of its own class position. In this paper, I suggest a way in which we can better understand the possible relationship of collective action and collective responsibility to a Marxist political project, broadly construed. I argue that the two discourses can have a positive impact on one another, and attempt to set the stage for a more in depth interaction between the two.

Where Social Ontology Meets Economic Anthropology: A Theory of Value of Intrinsically Worthless Things

Giuseppe Danese—Where Social Ontology Meets Economic Anthropology
Universidad del Rosario

Many ethnographic records exist of apparently worthless objects circulating in traditional societies— the most influential example being perhaps the Kula necklaces and armbands first described in Malinowski’s *Argonauts of the Western Pacific*. Marcel Mauss, among others, classifies these objects as money. The nature of money, and the way in which money and other seemingly worthless tokens acquire “value”, is a topic of great relevance to social ontology (cf. recent contributions by Lawson and Ingham in the *Cambridge Journal of Economics*). These contributions point to some fundamental flaws in the classical and neoclassical conceptions of “value.” To find instances of objects that seem to defeat these traditional account of value, one might look at more recent instances, discussed in the paper, such as Judy Garland’s “ruby slippers,” or home run balls that have given rise to an entire jurisprudence on lawful appropriators. This paper discusses these puzzling cases of worthless objects acquiring value within a social theory of value. The paper highlights significant overlaps between the social ontology approach of social totalities and social positioning, and the Maussian concept of value and exchange as a “fait social total.”

Groups As (Epistemic) Agents: A Conceptual Defence for Changing the Subject

Sven Delarivière—Groups As (Epistemic) Agents: A Conceptual Defence for Changing the Subject

Vrije Universiteit Brussel

The aim of this presentation is to contribute to a fruitful explicitation on the notion of an epistemic subject with a special focus on group understanding and its link to epistemic agency.

Traditionally, philosophers have taken for granted that the relevant epistemic subject (i.e. entity ascribable with understanding, knowledge,...) must always be individual humans. However, if we start from an ability-oriented conception of understanding and demarcate the understanding subject according to the system that implements those abilities then epistemic subjects can extend beyond, or be entirely different from, human individual. Allowing such unconventional subjects does require further justification because understanding is a cognitive ability which we want to ascribe to an epistemic agent. But what warrants being an epistemic agent? My contention is that such an agent is no more or less than a successful target of what I call the epistemic stance (very much in line with Dennett's intentional stance). The epistemic stance is successful if ascribing an entity with epistemic properties (e.g. beliefs, goals, problem-solving tactics) has explanatory or predictive power. This brings up the question of whether groups could be a useful target of the epistemic stance or whether we should keep our focus on its individual members only.

To explore the epistemic stance and its application to groups, I consider what I think are (a) the least and (b) the most convincing case of group understanding: From (a) a non-cooperating and random collection of people, in which case the group is no more than a shorthand for the aggregation of individual abilities, to (b) a complex dynamic of individual interaction that is isomorphic to the brain of an understander (e.g. Block's Chinese Nation), in which case the group-abilities can't be reduced to the abilities of its members (even though they are realised by them). Using these two extremes, I extract what I believe are the subject-changing factors that differentiates them: First, the degree of complexity in the group's working parts, which makes it desirable to bypass that complexity by referring to a macro-level instead (i.e. the group-level). Second, the possibility to discern the appropriate systematicities at this macro-level (i.e. epistemically relevant behaviour from the group), which makes employing an epistemic stance towards the group an efficient thing to do. And third, the macro-systematicities being emergent, which means the group properties are not reducible to the member properties. If the group is realized through the complex interaction between its members, there is no straightforward mapping-relation between the group macro-level and the member micro-level that implements it, which means those macro-systematicities are conceptually tied to that macro-level only, and one must change the subject to the group to talk about them.

Using these three factors, I then showcase their value in considering (more plausible) intermediate cases of groups and determine whether or not it is useful to change the subject to the group.

Social Categories and the Politics of Recognition

Robin Dembroff—Social Categories and the Politics of Recognition
Yale University

“Call me Caitlyn.” With these words, which overlaid Vanity Fair’s June 2015 cover photo, Caitlyn Jenner revealed her transgender identity to the world. But her words did more than reveal: they also demanded. Most obviously, “Call me Caitlyn” demanded that others recognize – i.e., acknowledge and affirm – Jenner as having a new identity: woman.

Reactions to Jenner’s demand were predictable. Jenner was warmly embraced and lauded by many on the political left for her decision to—as Jenner put it—live as her “authentic self”. Others rebuffed Jenner’s demands. Some conservatives did so savagely: blogger Matt Walsh called her a “mentally ill crossdresser”, and likened her gender identity to his 2-year-old son’s claims to be a Tyrannosaurus Rex.

Despite their differences, both reactions to Jenner manifested the same assumption: persons’ claimed identities should be recognized when and only when these are ‘true’ identities. Call this the insider assumption.

The insider assumption is dangerous. For one, it distracts from important questions like Who should be recognized as woman? with ontological questions about what makes someone a woman. For another, it frequently is used by white supremacist and religious conservative groups to demand respect for their identities, which entails respect for their opposition to (e.g.) gay, black, and women’s rights.

In this paper, I argue that a richer understanding of social categories and their relationship to recognition reveals that the insider assumption is mistaken. The categories underlying social identities – e.g., woman or disabled – often result from and perpetuate oppression. In such cases, the rules of access to these categories, or the social positions of their members, work to marginalize, exclude, exploit, or otherwise subordinate particular groups. Moreover, I argue, recognition is a tool that can be used to challenge or revise social categories. In light of this, recognition should not be constrained by the insider assumption: 1 instead, it should be used to shape, rather than merely reflect, the categories (and thereby, identities) operating within a social context.

My argument begins by outlining the insider assumption and its dangers. From here, I lay the ontological foundation for my argument by developing a model for oppressive categories, or social categories that arise from and perpetuate oppressive social structures. I then argue that recognition functions to either challenge or reinforce social categories. From here, it is a short step to see why the insider assumption fails: when social identities are based in oppressive categories, recognizing all and only the identities people already have will perpetuate the underlying categories. Given this, when deciding what identities to recognize, we should take more into account than what identities persons have. Rather, we must decide what identities we want there to be.

Unit-Of-Agency Dilemmas

Alexander Dietz—Unit-Of-Agency Dilemmas
University of Southern California

We often find ourselves in situations where we struggle to decide between two courses of action, each of which seems to be supported by compelling ethical considerations. Some philosophers claim that these situations can take an extreme form: that there are genuine ethical dilemmas, cases where some agent ought, all things considered, to do one action, and ought, all things considered, to do some other action, even though she cannot do both. This is a troubling prospect. As I will discuss, one of the key reasons why ethical dilemmas strike us as troubling is that it seems that they would make our struggle to decide what to do irresolvable. As Thomas E. Hill writes, in an apparent dilemma, conscientious people would find that "principles and values they assumed could never be compromised pull at them from opposite directions, threatening to tear apart that unity of soul long supposed to be the only indestructible reward of virtue."

Now, a number of philosophers believe that there are plausible ways to deny ethical dilemmas. But in this paper, I will argue that even if we are comfortable denying ethical dilemmas as they are traditionally understood, there is another kind of normative conflict that we must confront: cases where what ought to be done at one unit of agency is incompatible with what ought to be done at another unit of agency, or what I will call unit-of-agency dilemmas. In particular, several philosophers have recently argued that there are not only things that each of us individually ought to do, but also things that we collectively ought to do. However, if this is right, then there may be cases in which these obligations are incompatible. For example, there may be cases in which we should do something, but I shouldn't do my part. These cases, I will argue, threaten to be just as paralyzing as traditional "single-agent" dilemmas, but the problem cannot be avoided or resolved in the same ways.

The Reality of Institutional Kinds: A Causal Account

Zili Dong—*The Reality of Institutional Kinds: A Causal Account*
Simon Fraser University

This paper attempts to defend the reality of institutional kinds by developing a causal account of institutional kinds. According to the mind-dependence thesis, institutional kinds exist because they are collectively accepted; and it has been argued that accepting this thesis entails that institutional kinds are not real. However, I would argue that institutional kinds are causally (rather than constitutively) dependent on collective acceptance, which indicates that even though the mind-dependence thesis is true, institutional kinds are real.

First of all, I point out, all real kinds, including all natural kinds, are primarily causal kinds; a kind is real if and only if the members of this kind have certain unified causal characteristics which are robust enough to be taken as the essence of that kind. Given this criterion, many mental kinds and intentional action kinds (like vote and murder) are real kinds. Based on these results, I argue further that institutional kinds are causal kinds as well. My main argument for this is based on the causal theory of action: intentional actions are caused by mental states; especially, in order for my action to be intentional, the content of my mental states has to be causally relevant. Now suppose that we take an entity X as an instance of an institutional kind K (e.g., money) in a certain context; in this case, our collective acceptance of X should successfully cause certain institutional attitudes and actions (e.g., holding correct beliefs about the function of X; interacting with X in appropriate ways), which are necessary and sufficient for sustaining the existence of K. In other words, the dependence between the existence of K and collective acceptance of this kind is causal rather than constitutive. If this is true, the mind-dependence thesis will not be a threat to the reality of institutional kinds. A further implication is that it shows that as a kind of causal knowledge, our knowledge about institutional kinds is fallible; that is, collective acceptance may fail to create an institutional kind (e.g. common property) if the collective acceptance fails to be causally effective.

Specifically, I suggest that the causal characteristics of an institutional kind K may include at least the following aspects: (1) all members of K come from the same causal origin, namely, collective acceptance of K, as I have argued. (2) I would argue that K can be the node of many significant generalizations about causal relations involving K in the social world (by adopting a difference-making theory of causation). (3) Furthermore, K could have certain causal consequences which are also real kinds, as has already been argued by others.

Finally, I give a quick response to the following worry: my account heavily relies on higher-level causation involving social kinds; but there is no such thing. I argue that the difficulty raised by the causal exclusion argument will be no threat to higher-level causation if higher-level causation is understood in terms of a difference-making theory (as I have adopted).

Responsibility in Equilibrium

Hein Duijf—*Responsibility in Equilibrium*
Utrecht University

Suppose a bunch of drivers find themselves in the following situation on the US highway: all of them are speeding and each is aware that each is speeding over the legislated speed limit. The resulting pattern of behaviour is dangerous, the risk of imminent accidents is serious and could be diverted if each of them slows down. However, suppose that each knows, or expects, that the others maintain their current speed. Furthermore, assume that unilaterally slowing down will impose a huge risk of an accident, which exceeds the corresponding risk when all maintain their current speed. As a result, each driver maintains their speed and the dangerous pattern of behaviour is sustained. I raise two questions. Is anyone morally blameworthy for their involvement in the joint speeding? Does it make sense to say that the group of drivers is collectively blameworthy for their joint speeding?

The situation for two drivers can be represented in the game form of Figure 1, where each driver has two available options: to speed or to slow down. It is important that there is no time to communicate, they must simultaneously choose whether to maintain speed or slow down¹.

| | | Driver 2 | |
|----------|--------------|--------------|-------------|
| | | <i>speed</i> | <i>slow</i> |
| Driver 1 | <i>speed</i> | 1 | 0 |
| | <i>slow</i> | 0 | 2 |

Figure 1: The speeding drivers.

Does it make sense to say that the group of drivers is collectively blameworthy for their joint speeding? Although the answer need not be affirmative, I argue that the scenario is similar to institutional contexts and collective intentional action (Bicchieri, 2006; Guala and Hindriks, 2015; Guala, 2016; Gold and Sugden, 2007; Tuomela, 2005). The key reason for this similarity is that the speeding drivers are in equilibrium: each performs an action that is the best response to the others' actions; and each performs an action that best promotes her goal given her expectations. ² Moreover, it can be argued that the drivers collectively knowingly bring about an objectionable outcome.

Can any of the drivers be held individually morally responsible for their involvement in the joint speeding? Since the drivers are in equilibrium, this establishes that Driver 1 is individually blameless. After all, Driver 1 does the best she can given Driver 2's choice, and also in light of her expectations regarding Driver 2's choice.

I argue that something is missing from this standard individualistic analysis. After all, they were jointly able to avert the risks if all would slow down. This observation can be incorporated into the framework of team reasoning (Sugden, 2000; Bacharach, 2006; Gold and Sugden, 2007): individualistic reasoning recommends each to maintain speed, while team reasoning recommends

¹ Note that these numbers are meant to capture the moral value of the resulting pattern of behaviour. It should thus not be conflated with utilities, as commonly used in rational choice theory.

each to choose slow. ³ Hence, within the team-reasoning framework this gap is filled by the possibility of an agency transformation.

What Do Lisa, Her College’s Tango Club, and the U.S. Marine Corps Have in Common? Rethinking What It Takes to Be an Agent

*Michelle Dyke—Rethinking What It Takes To Be An Agent
NYU Philosophy*

In recent years, there has been growing interest in the notion of “group agency,” the idea that whole groups of human beings, such as wine clubs, budget committees, sports teams, churches, universities, or even international organizations, can count as genuine, unified agents in their own right. Yet is it true that a group like the US Marine Corps counts as an agent in the very same sense that a single human person, like Lisa, counts as an agent? I claim that, yes, they are both genuine agents. I offer an original characterization of the requirements of agency that, I argue, generates the right verdicts about what counts as an agent, especially when it comes to recognizing cases of genuine group agency. Other views do not generate the desired verdicts about what counts as an agent. One natural way of thinking about agents, generalizing from the everyday case of persons, is to assume that agents must be conscious beings with minds of their own. Yet the claim that agency requires a mind is obviously too restrictive if we are to make any sense of the claim that groups like wine clubs and sports teams can be genuine agents. Another way of characterizing the nature of agency stresses the idea that agents are rational beings. Yet this approach is wrong-headed because we want to make sense of the idea that it can be appropriate to criticize agents, including persons, for (even extreme) failures of rationality without casting doubt on whether they still count as agents. Other ways of characterizing the nature of agency, which have the appeal of seeming friendly to the possibility of group agents, have the problem of being too permissive about what counts as an agent. With the goal of defending the possibility of group agency explicitly in mind, List & Pettit (2011) give a characterization of agency that appeals to the possession of representational and motivational states (as well as a capacity for action) that are functionally analogous to human mental states. List & Pettit welcome the result that this characterization may count as an agent a simple robot designed to be able to realign cylindrical objects on a tabletop. In addition, the view would seem to count as agents a variety of other systems, like a computer algorithm designed to generate a list of prime numbers, or a heat-seeking missile. Yet intuitively, these things are not agents. My own characterization of the requirements of agency appeals crucially to the possession of a capacity for the self-determination of aims. I argue that entities that can determine their own aims are the sorts of things that we are right to hold responsible for their choices and behaviors, or to criticize for acting irrationally.

Were American Legal Realists Naïve Metaphysicians?

*Adam Dyrda and Paweł Banaś—Were American Legal Realists Naïve Metaphysicians?
Jagellonian University*

American legal realists (like Frank, Llewellyn, Cohen or Holmes, later as “ALR”) tend to be accused of many philosophical “sins”. E.g. they are perceived as philosophically naïve, eclectic and ametaphysical school of legal thought, skeptical of conceptual analysis as a useful tool of general jurisprudence. However, they perceived their conceptual and metaphysical skepticism as a virtue rather than vice.

In recent years the debate concerning the major epistemological and metaphysical presuppositions of ALR was somewhat revived. According to Leiter, ALR were methodological naturalists who tacitly employed hard-positivist concept of law as that concept was the one that figured in the most fruitful a posteriori research programs, i.e. the ones that give us the best account of how the world works. This suggests that ALR did not treat positivist concept of law as “real”, but at best – as pragmatically acceptable (and thus fallible). In such circumstances the question of the metaphysical status of ALR’s claims remains open.

In our paper we would like to analyze ARL’s main claims and argue that most of them can be plausibly interpreted as claims pertaining to so-called “naïve metaphysics” in the sense recently proposed by Kit Fine. According to Fine, naïve metaphysics do not, while fundamental metaphysics do – make essential use of the notion of reality and cognate notions. The questions concerning the former (the nature of things) can be effectively pursued independently from answering questions of the latter (whether these things are real).

We would illustrate the thesis that ALR were naïve metaphysicians by analyzing their metaphysical statements about legal rules, social rules, the nature of social change, the nature of law and legal indeterminacy. While ALR kept asking questions concerning the nature of legal facts, legal decisions and their “legal” grounds etc., they did not necessary trouble themselves with rising the primary question of fundamental metaphysics, namely the question whether they are real. So, they first asked questions what law is or may be like (or what it is or may be not) and did not bother to argue how it relates to the “concept of law” as denoting something real. They usually remained skeptical or indifferent about existence of rules or legal facts. Although they are traditionally called “legal realists”, they were interested in the structure of “legal appearances” rather than in the structure of “legal reality”. Eventually, they lost their faith in conceptual theories of law because they found most traditional answers to “fundamental” metaphysical questions about law hardly promising.

The naïve metaphysics of ALR did not, however, entail philosophical “quietism”. We would argue that apart from ALR’s commitment to “naïve metaphysics of law”, it did not preclude them from answering questions of “fundamental metaphysics of law”. Some realists (as our examples would show) took advantage of that possibility, but that did not necessary lead them to accept fundamental metaphysical claims of legal positivism or any other general theory of law. Metaphysical naïvety, not fundamentality, was an “essential” feature of ALR.

Truth and Social Ontology

Douglas Edwards—Truth and Social Ontology
Hamilton College

What is it that makes social and institutional claims like 'Andrew Cuomo is the Governor of New York', 'Angela Merkel is white', and 'Hillary Clinton is a woman' true? In this paper I develop an account of social and institutional truth in terms of the notion of superassertibility, or durable warrant to believe or assert a claim.

I begin by introducing the idea of superassertibility, and explain why it is a suitable candidate for understanding social and institutional truth. Superassertibility is a form of durable warrant to assert a sentence; a warrant which is available at some stage of inquiry, and then persists through all subsequent stages. The idea is defined by Wright as follows:

A statement is superassertible, then, if and only if it is, or can be, warranted and some warrant for it would survive arbitrarily close scrutiny of its pedigree and arbitrarily extensive increments to or other forms of improvement of our information. (Wright 1992: 48)

Why might this property be a suitable candidate for explaining social and institutional truth? We can note that it is generally taken to be a feature of the social and institutional domain that the existence of the objects and properties contained therein is dependent on what Thomasson (2003) and Searle (2010) call "collective acceptance". That is, the only reason there are governments, banks, races, and genders is because their existence is sustained by collective belief in, and acceptance of, such things. If no-one held any beliefs about presidents, or no-one accepted the legitimacy of the office of president, then there would be no presidents.

If we accept these ideas, then it tells us some things about how we should think about the truth of social and institutional claims. Firstly, it seems sensible to connect truth to what we have reason to believe, as the very existence of the entities that our beliefs and sentences are about is dependent on what we have reason to believe. That is, it does not make sense to think of truth as a relation between mind and a mind-independent reality, as we are stipulating that the reality in question is dependent on what we collectively believe and accept. Secondly, it gives support to the idea that, in the social and institutional domain, all truths are knowable. This is because of the dependence of the reality on belief and acceptance: if the existence of the objects and properties in question is dependent on beliefs, then it does not seem possible for there to be any truths that could extend beyond what we could, in principle, know.

I go on to consider some specific examples of social and institutional truths to apply the model. Firstly, I demonstrate how an account can be given of institutional truths, such as 'Andrew Cuomo is the Governor of New York'. I then show how social truths can be accounted for, using the examples of claims about race and gender as the key case studies.

Convention and the Nature of Institutions

Brian Epstein—*Convention and the Nature of Institutions*
Tufts University

Economists have long observed that social institutions matter for the development of nations. Acemoglu and Robinson (2012), for instance, examine macroscopic institutional structures, distinguishing the effects of “extractive” and “inclusive” institutions for economic growth. At a micro level, the prevailing economic theories of institutions and organizations treat them as rational solutions to strategic problems of interaction and coordination. In this paper, I argue against this micro approach to institutions, and in particular against the idea that institutions be understood as social conventions. If we are to construct a micro-level foundation of institutions to explain and design large-scale institutions, it will have to be different basis than the prevailing one.

To say that a social phenomenon is “conventional” sometimes means little more than that it is social, or that something about the phenomenon is arbitrary, or a matter of choice. In social theory, however, more precise analyses of convention are given, usually as variants of David Lewis’s 1969 analysis of convention, or else following analyses of convention that challenge Lewis’s approach. I propose a different approach to social convention, using tools from social ontology I introduce in *The Ant Trap* (2015). In the book, I develop the distinction between the grounding of social facts and the anchoring of frame principles. In this paper, I argue that conventions are best understood as frame principles. We can then distinguish several notions that are sometimes run together: (i) the convention (or frame principle) itself, (ii) the anchors or metaphysical explanation for the convention to be in place in a community, and (iii) the causes for those anchors to be in place. The difference among these can be seen by comparing a convention to a law (that is, another kind of frame principle): (i) there is a law, such as the definition and sanctions for insider trading, (ii) the facts in virtue of which that is a law, such as the enactment of the statute by the legislature and its interpretation by judges, and (iii) the causes for the legislature to have enacted that statute and for the judges to have performed the interpretations they have. Using these distinctions, we can characterize the conditions under which a frame principle is conventional: the conventional ones will end up only being a small subset—often an unimportant subset—of frame principles more generally. I further argue that it is misleading to analyze conventions in terms of structured attitudes or as strategic coordination devices.

Second, I consider the nature of institutions and their analysis in terms of strategic coordination. Prevailing theories of institutions regard them as made up of cognitive structures, such as sets of rules or attitudes, and they hold that institutions generally have the function of affecting cognitive structures, such as incentives, payoffs, and choices. I argue against both of these, describing more heterogeneous building blocks of institutions than just cognitive structures, and more general functions that institutions have than affecting such structures. And I argue against conventionality in particular. By clarifying the nature and functions of institutions, it becomes clear that only the rare instance is conventional.

How Do Social Groups Persist?

August Faller—How Do Social Groups Persist?

Cornell University

How do social groups persist? In this paper I build on Epstein's recent account in *The Ant Trap* to argue that social groups do not persist by enduring. The argument proceeds by applying to social groups classic puzzles of personal identity, the identity of ordinary objects, and vagueness. First, identity is transitive, but Parfit-style cases of groups splitting and fusing show that group identity is not. Second, identity is necessary, but Ship of Theseus-style cases show that group identity is not. Finally, identity is not vague, but Chisholm's Noah-Adam case and the Sorites paradox show that group identity is vague. Each case is particularly pressing for social groups.

In response, I consider solutions found in the literature on personal identity and coincident objects. Several are not plausible, as they do not apply to groups or do not address the puzzles. The space-time worm theory of Lewis is promising, but has negative metaphysical and semantic consequences, some new to groups. The best solution, I argue, is a close variant of Sider's: social groups are momentary stages, and group identity across time ought to be analyzed in terms of counterpart theory. I call this view "Group Stage Theory" (GST).

Next, I defend GST against objections. First, there is a worry that the argument over-generalizes to all grounded entities. It might, but the form of solution does as well. In particular, stage theory responds to Skiles' Ship of Theseus argument that grounding doesn't necessitate. Second, one might worry that on the stage theory we get the wrong number when counting groups over time. In response, I argue that our counting practices track groups "quotiented out" by a counterpart relation.

Finally, I illustrate how the view can contribute to certain politically-charged debates. GST provides a metaphysical and semantic framework for Haslanger's conceptual, descriptive, and analytical approaches to race and gender when describing groups across time. For example, consider the claim that the ancient Greeks were white. According to GST, this claim should then be translated as: the group constituted by the ancient Greeks (at some time t or across a range of times) is a group-counterpart of the group of white people now. Whether this is true depends on the counterpart relation in play. Approaches like Nell Irvin Painter's in *The History of White People* can be understood as using this flexibility in line with Haslanger's analytical approach, shifting the counterpart relation to highlight how a racial category came to exist under the social and economic conditions of colonialism. On this counterpart relation, Greeks were not white. The flip side of this flexibility is that other counterpart relations do deem the Greeks white. Even so, GST deflates the question to some dimension of similarity between the Greeks and contemporary white people.

Circularity in Searle's Social Ontology: With a Hegelian Reply

Jose Luis Fernandez—Circularity in Searle's Social Ontology: With a Hegelian Reply
Temple University

John Searle's theory of social ontology argues that there are indispensable normative components in the linguistic apparatuses termed status functions, collective recognition, and general agreement, all of which, he argues, make the social world. In this paper, I should like to examine Searle's argument for the emergence of social reality along two lines. First, I argue that Searle's ideas of status functions, collective recognition, and general agreement, which are the building blocks of social reality, are caught in an inescapable *petitio*. Toward this aim, I show how Searle's notion of status functions requires collective recognition; how his idea of collective recognition relies on general agreement; and how his idea of general agreement is dependent on status functions (hence the circle). Second, I show how ideas of status functions, collective recognition, and general acceptance have to presuppose culture in their functioning and existence. Here I note that Searle does not allow a space in which we can observe language in reciprocal relation to the institutions which not only are shaped by it, but which also shape language's practical applications. According to Searle, theorists that tried to show a connection between culture, language, and society, e.g., Weber, Bourdieu, and Habermas, all failed to see the constitutive role of language in the making of social reality. Subsequently, I argue that Searle is himself guilty of a certain kind of blind presumption, namely, that he takes for granted the role of culture (*Bildung*) in the creation and maintenance of language and human institutional reality. Toward this aim, I argue that Hegel's philosophy of culture avoids the *petitio* in Searle's theory by offering a more cohesive account of the reciprocal normative transactions between human beings and their social world.

An Ontological Analysis of Groups as Complex Variable Embodiments

Roberta Ferrario and Daniele Porello—An Ontological Analysis of Groups as Complex Variable Embodiments

Institute for Cognitive Sciences and Technologies and the Free University of Bolzano, Bozen

In this paper, we propose an analysis of groups and a first-order theory in which groups are viewed as social objects, in particular groups are constructed as variable embodiments (Fine, 1999) of structures of roles that are diachronically applied to collectivities of entities.

We assume an extensional mereology to talk about aggregates of objects (while still staying within first-order logic) and we show that a mereological aggregate of individual entities counts as a group whenever there exists a founding event that provides the aggregate of individuals with what we term its unity criterion. The unity criterion is understood as the description of the internal structure of the group, which is captured by the aggregate of roles that apply to individual members of the collectivity through time and that are generated by affiliation relations established between the members and the group.

To concede that a group may preserve its identity through time while possibly changing its members, we propose a notion of role which allows variable embodiments (by adopting Kit Fine's notion). We also loosen the constraint that groups may be formed only by human agents, including in our framework groups formed out of other groups or of artificial agents.

Moreover, we also concede that a group may preserve its identity through time, while possibly changing its unity criterion (i.e. the ascription of roles to the members or the structure of the group).

We propose then a few arguments for a multiplicative approach to groups that views them as new individuals – qua-objects à la Fine (1982) – which are in fact created when roles are enacted.

In particular, a multiplicative view allows us to predicate properties of groups, viewed as first-order citizens of the theory, to express identity statements about them and to reason about intentionality, agency and responsibility ascription.

On the one hand we formalise the dependence of the group on the founding event, on the collectivity of members, and on the unity criterion, by means of Epstein's notion of anchoring, on the other hand we use the notion of grounding to identify when intentions and actions can be legitimately ascribed to groups (Epstein, 2015).

Finally, we propose a taxonomy of groups, that separates gatherings (e.g. queues), informal groups (friends who meet on Friday to play volleyball), institutionalised groups (e.g. the PhD students of our program), group agents à la List and Pettit (2011), and organisations (e.g. corporations).

Intrinsic Masking and Sexual Orientation

Peter Finocchiaro—Intrinsic Masking and Sexual Orientation

University of Notre Dame

There is a longstanding dispute regarding the metaphysics of dispositions. Can a disposition be masked by features internal to the bearer of the disposition? In this paper, I argue that the metaphysics of sexual orientation provides a positive answer to this question. Yes, it is possible for a disposition to be intrinsically masked, and it is a relatively common phenomenon in our sex lives.

My argument turns on real-life cases about sexual orientation: gay Christians and political lesbians. In these cases, the individual is disposed to engage in a range of sexual behavior but, for deeply held moral reasons, decides to not engage in that behavior. I argue that, in these cases, the individual has a behavioral disposition that is masked by features intrinsic to that individual.

There are at least three reasons to care about intrinsically masked sexual dispositions. First, some have used the alleged impossibility of intrinsic masks as a device to define the distinction between dispositional and categorical properties. The fact that masking is sometimes intrinsic undermines this means of drawing the distinction. Second, reflecting on such cases helps to illuminate the correct manifesting conditions of sexual dispositions and, subsequently, helps to illuminate the nature of sexual orientation. Finally, the fact that there are intrinsic masks advances one line of reasoning regarding moral responsibility. My cases of intrinsically masked sexual dispositions bear a structural similarity to Frankfurt-style cases involving intrinsic interference -- for example, the case involving an addict who necessarily acts a morally blameworthy way. If individuals retain their sexual orientation in intrinsic masking cases, by parity of reasoning the addict retains the ability to act otherwise. Arguably, then, such cases do not count as counterexamples to the Principle of Alternative Possibilities.

In section one, I establish some background regarding the metaphysics of dispositions. Then, in section two, I argue for a theory of sexual orientation according to which sexual orientation is or is strongly tied to a cluster of dispositions regarding sexual behavior. In section three, I give two cases that strongly suggest that sexual orientation can be masked by features intrinsic to the individual. Finally, in section four, I sketch out some possible responses to my argument and show their deficiencies.

Ideal Theory and Shared Intentions

Christina Friedlaender—Ideal Theory and Shared Intentions
University of Memphis

In political philosophy, ideal theory is a methodological and justificatory framework for arriving at a set of normative principles (Rawls 1971; Robeyns 2008; Valentini 2009). Ideal theory emphasizes idealizations about agents and the societies in which agents exist and often presupposes a very particular picture of how social cooperation works. By removing the contingent features, or non-ideal conditions, of both individual agents and their social contexts, an ideal theory aims to capture the normative principles that ought to govern social cooperation. Such principles can then be extended to non-ideal social conditions with non-ideal agents through the development of a non-ideal theory (Phillips 1985; Stemplowska 2008). However, these normative projects about social cooperation rests on descriptive projects about the terms of social cooperation. In *Theory of Justice* (1971), for instance, Rawls' normative project rests on descriptive claims about how individual planning agency functions in relation to his model for joint deliberation about shared ends.

I argue that the methodological framework of ideal theory has been taken up in standard accounts of shared intention. In many standard accounts of shared intention, there is an emphasis on developing success conditions that begin from idealizations about human agents and exclude non-ideal conditions in which agents develop and act, individually or jointly. In an attempt to offer universal success conditions for shared intention, many of these projects offer asocial and ahistorical accounts of human agency under the assumption that so-called contingent features of agency or non-ideal conditions of agents' social contexts can be filled in later. Focusing on the work of Michael Bratman and Margaret Gilbert, I show how their projects rely on the methodology and justificatory structure of ideal theory.

Broadly, I capture how ideal theory functions in both political philosophy and shared intentions in order to set the foundations for critique. In the political case, ideal theory aims at normative principles about social cooperation by first offering descriptive claims about social cooperation. As Charles Mills argues in *"Ideal Theory" as Ideology* (2005), these descriptive claims are already normatively-loaded insofar as the idealizations used in ideal theory do not show what an agent is, but what an agent should be. For Mills, the normative prescriptions offered by ideal theory cannot extend to the non-ideal world because such principles are derived from the ideal world. Moreover, such normative prescriptions are derived from a normatively-loaded description of agency. The use of ideal theory in shared intentions has determined both normatively-loaded descriptive claims about how two or more agents share an intention as well as explicit normative claims about sharing an intention. My project shows how ideal theory underpins standard accounts of shared intention. By establishing this connection, I argue that if the use of ideal theory is concerning in the political case, as Mills takes it to be, then it is also of concern in the case of shared intentions.

Shared Intentions and Intending to Merely Do One's Task

*Caitlin Fitchett—Shared Intentions and Intending to Merely Do One's Task
Harvard University*

Michael Bratman reductively explains the intentions of group agents in terms of the intentions of individuals. On his view, a group agent's intention to J is explained by the individuals who comprise that group each holding an intention of the form "I intend that we J." In holding such a reductive account of group intentions, Bratman is able to give a ready response to questions about group membership: what it means to be a member of an intentionally acting group is to share an intention with the other members of the group. I argue that Bratman's view is challenged by cases where individuals who appear to be members of a group agent deny that they share an intention with the other apparent group members.

Consider the following case:

Factory Case: A factory is producing chemical weapons. The production of these weapons is intentional. When each factory worker is asked what he or she is doing, each responds, "I am merely completing my task." And, when the workers are asked what they intend to bring about, each replies, "I intend to complete my task, and I understand that in completing my task I help to make the chemical weapons that the factory produces. But I don't intend for these weapons to be produced. I merely intend to do my own, very specific task in order to get paid."

Because each factory worker does not intend to produce chemical weapons there is no shared intention to produce the weapons. But given this, Bratman is unable to explain how it is that the factory intentionally produces chemical weapons. I motivate the idea that the case does serve as a counter-example to Bratman's theory, and in doing so also argue that we should view the factory as being a group agent. Following this, I argue that the workers do not need to intend to produce the chemical weapons in virtue of foreseeing that their actions will lead to the making of the weapons. Nor do the individuals need to intend their tasks under a description that specifically references the production of the weapons. In making these arguments, I block two routes that might be adopted by a supporter of Bratman's view in order to argue that the factory workers are irrational and that if they intending rationally, then there would be a shared intention.

Intuitions About Group Minds: Public Evidence

Brian Flanagan—*Intuitions About Group Minds: Public Evidence*
National University of Ireland, Maynooth

There are two basic approaches to the possibility of group action in circumstances of inter-member disagreement, i.e., to the decisions of collegiate courts, boards and cabinets. Whereas individualists reject the possibility of group agents (e.g., Weber 1922; Hayek 1942; Elster 1985), collectivists acknowledge their reality. List and Pettit 2011 distinguishes ‘thick’ collectivism, which holds that groups have minds of their own (e.g., Tollefsen 2002; List and Pettit 2011) from ‘thin’ collectivism. Thin collectivists recognize that, as a group member, one may make a decision or hold a belief that conflicts with one’s own ordinary individual attitudes. But they insist that any such group decision or belief will derive systematically from the minds of the group’s members alone (e.g., Copp 1995; Kornhauser and Sager 2004). I show that, in respect of both the question of the proper application of Occam’s razor principle and of the nature of collective moral responsibility, the debate between thin and thick collectivists is very nearly as consequential as that between individualists and thick collectivists. A recent argument for thick collectivism due to Christian List and Philip Pettit has gained wide influence.

All varieties of collectivism suppose the intuitiveness of group action, i.e., that ‘[c]ommon sense... represent[s] many collections of human beings as if they were unitary agents’ (List and Pettit 2011, 1). To ground thick collectivism, however, the assumption that our commitment to the truth of group action sentences can be redeemed only by the extension of agency to groups is insufficient. As against the thin collectivist, thus, List and Pettit construct an argument that assumes that our commitment to the truth of group action sentences is constrained by the application to groups of norms of rationality but unconstrained by the imperfect capacity of rules of attitude aggregation to satisfy those norms. I show that, to accept List and Pettit’s argument, we must be able to identify at least one example of the ascription to a group of a decision contrary to the non-rational indication of the relevant rule of attitude aggregation.

Recreating circumstances conducive to the advertence to group minds in counter-factual contexts presents difficulties. Crucially, factual contexts derived from the public record may feature consideration of the ascription of group action by those with prior history of engagement with the relevant group. Moreover, it is part of the common law tradition of legal adjudication that the members of multi-members courts of law publish not only their respective views on disposition of litigation but also on the reasoning underlying those views. This feature produces factual contexts in which the relevant rule of attitude aggregation indicates a non-rational court decision. Accordingly, having first presented a formal reconstruction of List and Pettit’s argument for thick collectivism, I test its assumption that the ascription of a group act is conditional on norms of individual rationality by reference to the historical ascription of group acts of legal adjudication. We find that, on examination, intuitions about group agency offer no reason to posit minded group agency.

Can Gender Be Part of the Plan?

Jade Fletcher—Can Gender Be Part of the Plan?

University of Glasgow

The Canberra Plan endures as one of the most promising approaches to philosophical theorising. It promises an exhaustive theoretical enterprise. It aims to synthesise a fairly traditional conception of philosophical analysis with metaphysical naturalism. The project starts from our pre-philosophical beliefs about the world. These platitudes of common opinion are systematised and subjected to conceptual analysis, thus deriving truth conditions for the relevant sentences. Finally, any worldly satisfiers of those truth conditions must be identified. One putative advantage of the approach is that it respects our pre-philosophical beliefs, such that we get to take much of our ordinary talk as true, whilst incurring minimal ontological inflation.

In this paper I argue that the Canberra Plan is ill-equipped to offer a satisfactory metaphysics of gender. Insofar as the Canberra Plan aims to provide a general and unified approach to philosophical theorising, this is a significant problem. I argue that this deficit in their method stems from the robust role assigned to pre-theoretical beliefs in constructing philosophical analyses. This argument employs a Marxian conception of ideology to explain why our pre-theoretic beliefs about certain social kinds are likely to deliver politically dubious metaphysics of the social world. The first half of the paper is dedicated to exercising this theoretical shortcoming. In the second half, I suggest a way in which the Canberra Plan can address and rectify this problem, with a view to maintaining the theoretical viability of the Canberra Plan with respect to such politically important concepts.

I proceed as follows. First, I outline David Lewis's project of philosophical analysis, highlighting the integral role played by common opinion in this project (§I). The Canberra Plan methodology is associated with many philosophers; I therefore focus my discussion on Lewis as he offers the most comprehensive and systematic approach. Second, I show that, as it stands, the method offered by Lewis cannot satisfy the needs of a politically informed analysis of gender concepts (§II). Third, I utilise Sally Haslanger's ameliorative analysis of gender concepts, as a means of illustrating an approach to a politically informed and useful metaphysics of gender (§III). Third, I offer a positive proposal for how the Canberra Planner can adapt their method, whilst incurring minimal costs, by analogising Lewis's treatment of the tension between materialism and common opinion (§IV). Finally, I motivate my claim that, like materialism, social justice can be a constraint on the uptake of common opinion (§V). I articulate two kinds of arguments in favour of my proposal. First, a non-normative argument, according to which materialism and social justice can act as constraints as that are both guides to truth. Second, a normative argument, according to which materialism and social justice can act as constraints as they both facilitate the explanatory purposes of their respective domains.

According to Others

Carmel Forde—According to Others

Saint Mary's University

Although I am sympathetic to Lynne Rudder Baker's stated goals, in her argument for inclusion of persons in social ontology, I believe that her account of human persons may be enriched through an examination of the insights of Maurice Merleau-Ponty on the body-as-subject, and Lev Vygotsky's writing on the impact of self-constitution of the social speech of the child.

In part my goal is to delineate the meaning and significance of the term "intentionality" by examining Merleau-Ponty's analyses. There are two common meanings of the term, found in the literature on social ontology. The frequent reference to "goal-directedness" in the deployment of such terms as "we-intending", and "shared intentionality" is one. This use however seems to thwart the promise of intentionality (understood as a directedness toward meaning) for a richer account of "shared intentionality". In part, the phenomenological analysis of intentionality provides a reprieve from the focus on representationalism, which itself is problematic for some issues in social ontology.

I reach into Merleau-Ponty's corpus of writings, particularly on the infant and child, for an appreciation for the phenomenological meaning of intentionality. I believe his accounts can provide us with an account of the nature of collective intentionality, wherein we can meaningfully discuss shared beliefs about the world, or shared attention to the world, which is fundamental for any collective agency in transforming the world. I focus especially on his discussion of the body-as-subject, and on the body's intentionality. When we speak about human beings as intentionally related to meaning we identify them as perceivers and agents that share a world, whose lived experiences of that world overlap substantively, and who may also share projects and goals. By providing a detailed account of the significance of the phenomenological account of intentionality through some key examples of its pre-reflective use, social ontologists might find help in articulating meaning for "shared intentionality" "believing jointly", how to analyze social objects, and other issues in social cognition.

I develop my argument by examining Merleau-Ponty's analyses of motor intentionality, specifically on imitation of adults by infants, and on distinctions he makes between abstract and concrete movement, and the implications of these examples for shared intentionality. I further examine Lev Vygotsky's insights into the development of inner speech, and analyze the example of silent reading to present a case for considering pre-reflective intentional relations to meaning as important for socially constituted facts, institutions, and other meaningful referents in social ontology.

An Argument Against Causal Explanations of Akrasia

Megan Fritts—An Argument Against Causal Explanations of Akrasia
University of Wisconsin-Madison

Causal theories of action, like that of Donald Davidson, are commonly regarded as the “standard view” in philosophy of action. However, causal explanations of the paradoxes of akrasia—that is, of how someone could judge one option best without desiring it, or how someone could knowingly choose to undermine their own goals—encounter an unexpected difficulty. In this paper, I argue that if there is such a thing as group agency (and it is plausible that there is), then causal theories of action cannot explain group akrasia—a failure that has lessons for how we understand agency and individual akrasia, as well.

I begin by discussing the nature of group agents. Although my ultimate argument is conditional—if there are group agents, then causal explanations won’t work—I spend some time detailing the terrain of arguments for realism about group agents. Next, I discuss what it means for a group to have attitudes, such as judgments, beliefs, and intentions, and show why attempts at reducing group attitudes to the attitudes of the individual members fail. I then argue that attitudes of groups can be identified in the same ways that the attitudes of individuals are identified – through assertions, dispositions, activity.

Next, applying this conception of group agents, I detail several examples of group akrasia (group agents which act contrary to their judgment about what they ought to do), arguing that the beliefs and desires of neither the individual agents nor the group agents could have caused the resulting akratic action. Therefore, either the causal theorist is mistaken about all causes of akrasia, or akrasia is multiply-realizable. If akrasia is multiply-realizable, then the causal explanations of akrasia fail because their explanation is a bad one. But, in virtue of what are causal explanations bad? Good explanations, I contend, carve at natural joints. In my paper I compare the causal explanations of akrasia to attempts at theory reduction and bridging laws in science, where these reductive explanations failed to capture all cases of the relevant phenomena. I argue that the explanations of akrasia which fail to explain some instances of akrasia are missing something important about the subject matter.

I consider two different objections to my argument. The first objection argues that, because the subjective experience of temptation is essential to akrasia, and group agents are not subjects, they cannot be akratic. I respond that this objection conflates subjects and agents; insofar as they are conceptually distinct, we ought not to hold one to the standards of the other. The second objection argues that, if part of the function of beliefs and desires is to cause action, then the relevant beliefs and desires can be identified wherever there is an action taking place. I respond by arguing that identifying group beliefs and desires ostensibly ignores other evidence about the group’s attitudes, and results in uninterestingly circular definitions of attitudes and actions.

Against the Collective Responsibility of Groups

Christopher Fruge—Against the Collective Responsibility of Groups
Rutgers University

In this paper, I provide three arguments that groups do not have the sort of freedom required for moral responsibility. These arguments are modeled on three different incompatibilist arguments for the conclusion that individuals do not have the sort of freedom required for moral responsibility.

Much of the debate over individual freedom has focused on whether responsibility is undermined by features of causation or diachronic determinism. The two most prominent incompatibilist arguments from diachronic determinism are the consequence argument and the manipulation argument. One version of the consequence argument holds that given diachronic determinism the past and the laws of nature necessitate any future event, and so for any given action there are not the relevant sort of alternative possibilities required for responsibility. The manipulation argument holds that agents who perform actions which are diachronically determined cannot be held to be responsible in the same way that agents who are manipulated cannot be held to be responsible. The most prominent argument from causal considerations is the exclusion argument against the causal efficacy of action. The exclusion argument claims that if actions are not reduced to low-level physical events, then they have no causal power, which is a requirement on moral responsibility.

Individuals are not the only entities that act. Groups act as well. But their actions are metaphysically determined by the actions of their members and the relations between them. Such metaphysical determination is analogous to causation: the interrelated actions of the members bring about the actions of the group. Moreover, the interrelated actions of the members synchronically necessitates the activities of the group in a manner akin to diachronic determinism: given the interrelated activities of the members and the principles of how they give rise to the group's activities, the group's activities are necessitated.

In light of the parallels between causation and diachronic determinism on the one hand and metaphysical determination and synchronic necessitation on the other, I construct three arguments for the conclusion that groups are not responsible for their actions. These are modeled, respectively, on the three different incompatibilist arguments against individual responsibility. First, I offer a consequence argument that the synchronic necessitation of the groups' activities by the interrelated actions of its members precludes the sort of alternative group-level actions that are required for responsibility. Second, I offer a manipulation argument that such synchronic necessitation undermines responsibility in the same way that manipulation of the group undermines responsibility. Third, I provide an exclusion argument to the effect that unless the actions of the group are reduced to the actions of its members, then they do not have the sort of causal power required for responsibility. However, if they are so reduced, then the responsibility properly lies with the individual members, not the group itself. Taken together, these three arguments provide a strong case against the collective responsibility of groups.

Joint Intentionality and Embodiment: A Performative Account of Collective Action

Barbara Fultner—Joint Intentionality and Embodiment: A Performative Account of Collective Action

Denison University

What does it take for two or more people to act together in political contexts? What is the relationship between joint agency and solidarity? Many, if not most, social theorists have thought that solidarity requires unity of purpose, perhaps even of identity. But as feminists and others have increasingly drawn our attention to the diversity among groups making up social movements, such unity seems harder and harder to come by. What are the conditions of possibility of joint agency across difference? On the one hand, recent research in joint agency and social ontology (e.g., Searle 2010, Gilbert 2013, Bratman 2014, Tuomela 2013, Epstein 2015, Ludwig 2017) has largely failed to tap feminist scholarship and critical theory as a resource. Paying attention to the social construction of identity of individuals as well as groups, the relationality of the self, and the role of power fundamentally reframes the problem of collective intentionality. On the other hand, feminists have placed solidarity front and center of their theoretical and political projects (e.g. Mohanty 2003, Dean 1996, Weir 2013) and routinely refer to collective bodies and collective action, yet detailed analyses of joint action are largely lacking in feminist discourse. Rather, they have tended to conceptualize solidarity in terms of identity rather than joint action. (Hannah Arendt's definition of solidarity as action in concert is an important exception here.) The literature on collective agency may thus be helpful for feminist theory.

In this talk, I bring into dialogue two unlikely conversation partners, Judith Butler and Raimo Tuomela, in order to develop a view of joint action that does justice to its embodied nature, to the multiple forms it can take, and the multiple modalities it involves. I focus on Tuomela's fine-grained categorization of collective action that differentiates between different modes of acting in concert as a resource for feminist conceptions of collective action and solidarity. Nonetheless, any such categorization must be examined and reframed in light of the above feminist critical insights. In particular, I argue that Butler's emphasis on performativity and the bodily component of acting in concert is vital for an account of collective performativity. Although her focus on large public assemblies and movements like Black Lives Matter leaves out an arguably more basic joint agency at the level of dyads and small groups that greatly helps us to understand collective agency in general, collective performativity nonetheless offers a rich resource for exploring new forms of solidary sociality.

Rationality and the Acquisition of Intentions: Why Rational Choice Theory Is Not Self- Defeating

Roberto Fumagalli—Rationality and the Acquisition of Intentions: Why Rational Choice Theory Is Not Self-Defeating
King's College, London

The critics of rational choice theory (henceforth, RCT) frequently claim that RCT is self-defeating in the sense that agents who abide by RCT's prescriptions are less successful in satisfying their preferences than they would be if they abided by some theory of rational choice other than RCT (e.g. Bratman, 1987, Gauthier, 1997, Kavka, 1983, McClennen, 1990). In this paper, I combine insights from philosophy of action (e.g. Audi, 1991, Mele, 2000), philosophy of the cognitive sciences (e.g. Bovens, 1995, Gold and Sugden, 2007), and the normative foundations of RCT (e.g. Hausman, 2000, Sugden, 1991) to rebut this often-made criticism. I then explicate the implications of my thesis for the wider philosophical debate concerning the normativity of RCT for both ideal rational agents and real-life agents having limited control over the formation and the dynamics of their own intentions.

The paper is organized as follows. In Section 2, I draw on some putative paradoxes of RCT (e.g. Kavka's, 1983, toxin puzzle) to examine the issue whether an agent who abides by RCT's prescriptions can form so-called degenerate intentions, i.e. intentions to perform actions that maximize the agent's lifelong payoffs while failing to maximize her subsequent payoffs. I then argue that, contrary to what these putative paradoxes are often taken to indicate, agents can rationally form degenerate intentions, and therefore such putative paradoxes fail to show that RCT is self-defeating. In Sections 3-6, I defend my thesis against four major objections put forward to demonstrate that RCT is self-defeating: the objection from psycho-physical inability (e.g. Gauthier, 1998); the objection from temporal situatedness (e.g. Bratman, 1999); the objection from bootstrapping (e.g. Clarke, 2008); and the overdemandingness objection (e.g. Holton, 2009). If my thesis is correct, prominent attempts to demonstrate that RCT is self-defeating do not withstand scrutiny, and RCT has much greater normative tenability than its critics allege.

Analyzing Social Construction, Dismantling Wide Constructionism

Jorge Garcia—Analyzing Social Construction, Dismantling Wide Constructionism
Boston College

After considering whether the language of construction should be taken more literally or less, and given reasons for the former option, I undertake to analyze social construction by examining what we should understand to be its agents (who can and does construct), process (what they do that amounts to constructing), preconditions (in what circumstances their performing those actions can constitute constructing), raw materials (out of what they construct), tools by using what, and in which ways, they construct), and products (what can be constructed).

I compare and contrast the elements I identify with similar aspects of Searle's large-scale conception of social construction, which he has elaborated and revised over several books. I address three important questions that his work poses and attempts to answer, criticizing his own responses to them and suggesting improvements.

Further, I contrast my analysis of construction with alternative accounts sketched in the literature, especially, by Hacking and Haslanger, and conclude by drawing lessons from my account about the narrow limits on what we can plausibly be claimed to construct, and offer a reminder that nothing that is necessarily internal to construction—including persons, their mental activities, time, places, agreement, communication—can itself have been originally constructed.

To adapt Pettit's language in discussing Hobbes, in my account, only things of a certain type can be made from, of, and more important, by using, words.

On the Possibility of Group Forgiveness

*Benjamin Genta—On the Possibility of Group Forgiveness
Munich Center for Mathematical Philosophy*

Forgiveness is an essential act in human relations. In Hannah Arendt's words, if it weren't for forgiveness we would all "be confined to one single deed from which we would never recover" (1958). With the recent rise of truth and reconciliation commissions in countries plagued by mass atrocities, and the ever more frequent scene of politicians asking their constituents for forgiveness, the following question arises: can groups forgive? Although some philosophers have argued for a variant of group forgiveness (see Digeser 2001, Wolterstorff 2013, MachLachlan 2012), I do not think any has done this sufficiently.

In this paper, I argue that a specific type of group can be said to forgive, although perhaps with some limitations compared to the interpersonal case. I do so by first briefly presenting the generally accepted account of interpersonal forgiveness in the literature, which takes forgiveness as being a voluntary forswearing of resentment.

I then show that there are some important shortcomings with this account, and provide what I argue is a more-encompassing definition; namely, I define forgiveness as a voluntary process where one commits to set up one's attitude to overcome one's resentment after a wrongdoing. This definition relies on the notion of commitments as an uncertainty-reducing tool (Pacherie 2015). I show that this new definition still satisfies Hieronymi's (2001) three criteria for interpersonal forgiveness, and makes precise certain subtleties of the accepted view.

Then, per French's distinction (1984), I give a brief account of three types of groups: conglomerates (e.g. teams, social clubs); non-random aggregates (e.g. African-Americans, the Jewish community); and random aggregates (e.g. people at Central Park at 11:00). I proceed to briefly show the failures of previous accounts of group forgiveness.

Finally, I argue that conglomerate groups can forgive another agent. I show that this is the case for two different notions of collective action: Bratman's shared agency (1999) and Gilbert's joint-commitments (2013). Under both accounts, group forgiveness still satisfies Hieronymi's three conditions. Although, the use of commitments in our new definition of forgiveness should solve some conceptual difficulties that might arise in thinking of group forgiveness, I will note some differences between the two.

I end the paper by providing a roadmap for future research in the topic of group forgiveness and its possibility. Some open questions are: can non-random aggregates forgive? Are groups ever required to forgive?

My hope is that this paper will provide a resourceful step in expanding the discussion of collective forgiveness in philosophy, which is not only philosophically interesting, but could be favorable for post-conflict reconciliation and ameliorating relations between groups traditionally at odds with each other.

Deconstruction of Gender, Reification Of Race: Inconsistent Theories of Social Ontology on the Left

Samantha Godwin—Deconstruction of Gender, Reification Of Race: Inconsistent Theories of Social Ontology on the Left
Yale University

Feminist philosophy since at least the second wave, has historically interpreted gender as the contingent social constructs built up around, and correlating to, biological sex. Race has long been recognized as a set of socially constructed categories that vary according to time and place, lack any discrete biological reality, and have historically served to entrench social divisions. While contemporary feminism still nominally endorses the social constructiveness of both gender and race, feminist and left political thought has adopted new, incompletely articulated social ontologies of gender and race. Gender, once understood as the social ascriptions imposed according to sex to support patriarchal power relations (MacKinnon, Firestone, Haslanger) is now reimagined as a matter of liberating individual expression and personal self-identification. Race, once likewise understood as a social invention to ideologically justify caste hierarchies, slavery, and intra-class division (Marx), has been reendowed by the left as a unchanging, inherited trait that must be statically preserved and forms the basis for asserting group rights and collective guilt. Regarding gender, feminist and leftwing thought has adopted a hyper-individualism that transcends social power relations, but regarding race, feminist and leftwing thought has revived a reified, essentialist kind of nationalism that subsumes individual difference into monolithic group holism.

These two implied social ontologies of phenomena previously considered similar could not be further opposed to each other or less compatible. Yet rather than examining the profound tension between contemporary feminist conceptions of gender and race, these tensions have been largely neglected, or outright suppressed, as most vividly witnessed in the “Hypatia affair.”

These diverging implied social ontologies encourage difficult to reconcile couplets of political prescriptions regarding race and gender. Gender is believed to be fluid and a legitimate subject of contestation, self-definition, and culturally subversive queering and drag; but race is fixed and cannot be legitimately contested or self defined – to even adopt, let alone subvert, the cultural practices of another race is thought impermissible appropriation. Transgender identification is thought to deserve dignity and respect, but transracial identity scorn and ridicule. Gender neutrality is thought to be a contemporary moral imperative, but race blindness is an anachronistic liberal vice. To consider biological sex as giving rise to differing political interests is to adopt an intolerable “TERF” bigotry, but there is nothing segregationist or problematic about self-organization into black-only spaces. While intersectionality theory in its popularized incarnation dispels the notion of any shared experience common to women (MacKinnon, Snyder), the conflation of disparate populations into a common ‘person of color’ category is thought to promote the right kind of unity. Third wave feminism has torn down grand narratives about patriarchy and gendered power relations (Snyder), while simultaneously promoting grand explanatory narratives concerning white supremacy and anti-blackness.

These contradictions in feminist thought on race and gender are an obstacle to advancing

egalitarianism on an moral universalist basis. They both render the inequalities attendant to sex qua sex impossible to recognize, while also subsuming the material impact of racism on individuals into narratives of races as monolithic groups.

Harmonic Ontologies

Matthew Goldstein—Harmonic Ontologies
Dartmouth College

In his essay “On What There Is”, W.V.O. Quine holds that a participant in a given discourse commits herself to an ontology including all and only those objects over which the discourse's bound variables—words like 'something' and 'everything'—must be able to range in order to render the discourse's assertions true. He also argues, in “Identity, Ostension, and Hypostasis”, that these objects may include in their ranks spatio-temporally disparate summations. For example, a river is a summation composed of all its many momentary river-stages that come and go as water flows downhill.

In this paper, I apply these ideas and others from Quine's philosophy to the questions of whether and how we commit ourselves to ontologies that include music. We may listen to Beethoven's Symphony No. 7, play it, and read its sheet music. It is unclear, however, what exactly we hold exists when we engage in discourse about it. Do we merely stipulate the existence of the vibrations that convey sound to a listener, or the existence of the symphony's sheet music itself? Intuition would likely incline us to say no. This paper proposes, following Quine, that Beethoven's Seventh and other pieces of music are rather meta-summations of spatio-temporally disparate objects, somewhat like rivers. Further, any discourse in which we speak of Beethoven's Seventh simpliciter is one in which we are more concerned with what unifies all its renditions than what distinguishes one from another. Therefore, it is appropriate under Quine's proposed maxim of the identification of indiscernibles to speak of the symphony as a unified object. Finally, because in such a discourse Beethoven's Seventh is treated as a distinct whole, the bound variables of that discourse must range over it to render the discourse's assertions true. This means participants in that discourse commit themselves to an ontology in which the symphony exists.

This paper is concerned with an investigation into the linguistic patterns that allow us to incorporate even unusual and counterintuitive “objects” into our chosen ontologies. Because of the variability involved in when, where, and how a piece of music is performed or presented—in record, sheet music, or auditory form—it bears resemblance to a great many slippery social phenomena. Quine's tools allow us to understand how we may commit ourselves to harmonic ontologies in which pieces of music exist. They may therefore also provide us the groundwork for understanding our ontological commitments to these phenomena.

Racial Objects: A Fanonian Account of Race

David Gray—*Racial Objects: A Fanonian Account of Race*

University of Memphis

Having a concept of race seems crucial to contemporary social and political discourse. Race seems to be a significant factor for employment, quality of education, likelihood of incarceration, health, and other determinants of our life prospects. This all seems to imply the existence of races; and, if races exist, we should be able to explain what they are.

Much contemporary work in racial ontology proceeds by using the standard philosophical tools of conceptual analysis and/or reflective equilibrium. We have seen how such tools, even in a limited context such as the U.S., fails to yield consensus among 'ordinary' folk (Glasgow et al. 2009). Secondly, 'specialist' accounts of race that appeal to scientific or sociological experts share the dubious assumption that ordinary folk can and do defer to them for the meaning of racial terms. I say this is a dubious assumption because there is no expert consensus of what races are, and thus no designated experts to which we may defer.

In contrast to the above methodologies, I propose what we might call a 'racialization first' approach in which we determine what must be minimally true of how races are conceived given the existence of racialized practices and the formation of racial identities (in virtue of racialized treatment). This account is ordinary: even though it does not inquire into how ordinary individuals think of races (avoiding the first problem mentioned concerning conceptual analysis) it examines ordinary racialized acts. Additionally, by exploring ordinary racialized acts, this approach avoids the second problem mentioned concerning expertise and deference.

This approach requires an examination of astute observations of racialized. This approach is inherently piecemeal, but starting with Fanon, I will argue, proves to be highly insightful. While Fanon does not have a theory of race, per se, his account of the effects of racialized treatment on racial identify formation is among the most insightful we have and is well suited for use in developing a theory of race. Fanon's experiences in Martinique and Algeria are inseparable from his discussions of racialized treatment and identity formation, and can deeply inform what I take to be global features of our understanding of race.

Using Fanon's account of racialized treatment and the formation of racial identities, I will develop an account of race that has both descriptive and normative aspects. Races, I will claim, will fall into a category I call 'natural artifacts'. Natural artifacts are composed of a group of natural properties aggregated together to serve a functional role. Thus, races can be thought of as the aggregation of a group of natural properties, shared by a set of individuals, that can serve a functional role (e.g. economic gain). Additionally, insofar as these natural artifacts are objects, there is an inextricable normative dimension to this concept of race: persons are viewed in such a way that the aspect of humanity that would entitle them to not being treated solely as means to an end (viz. agency) is either denied or ignored. The group of properties that make someone a racial object are smaller (or at least different) than those that make someone a person.

Anchoring Human Social Kinds

Aaron Griffith—*Anchoring Human Social Kinds*
College of William & Mary

Epstein draws an influential distinction between 'grounds' and 'anchors' in his (2014, 2015, 2016). The grounds g for a social fact f are the facts that provide the metaphysical reason for f 's obtaining. Anchors, on the other hand, fix or put in place the conditions under which facts of one sort (G 's) ground facts of another sort (F 's). Much of the discussion of anchoring concerns examples of social kinds we explicitly recognize as such, e.g., money and US Senator. In many of these cases, anchoring their grounding conditions is largely a matter of the collective acceptance of those conditions. Human social kinds such as race and gender, however, do not seem to fit these paradigms. These kinds are widely believed to be natural or non-social kinds despite the fact that (according to many) they are social kinds. If these kinds are not widely recognized for what they really are (viz. social kinds), then anchoring the actual conditions under which they are instantiated is not a matter of collective acceptance of those conditions.

The aim of this paper is to give an account of the anchors for the conditions under which a person belongs to a human social kind like race and gender. More precisely, the account identifies the facts that anchor conditions of the following form: For all persons P , the fact that P has instances of properties G_1, \dots, G_n grounds the fact that P is a member of human social kind K . The paper begins by demonstrating the inability of accounts that appeal only to collective acceptance to identify the correct anchors for such grounding conditions. Then the paper offers an ontology of kinds like race and gender, one that will make the identification of what anchors their grounding conditions tractable. Developing views found in Haslanger (2012; 2016), Ritchie (ms.), and Mills (1997; 1998), races and genders are treated as positions embedded in social structures, where these positions are individuated by the function or role their occupants play in the structure. These roles are characterized by various norms and powers that are applied to/bestowed on the occupants of these roles, e.g., norms of behavior and treatment, social abilities, and opportunities. On this view, individual persons belong to a race or gender by instantiating certain properties by which they play the role characteristic of the race or gender.

This framework, it is argued, makes perspicuous what facts anchor the grounding conditions for human social kinds. What puts in place such conditions, to be explicit, are facts about the relevant social structure, its positions, and the various properties relevant to performing the functions characteristic of the positions. The account may be represented in the following way, inspired by Epstein's discussion of various anchoring schemas (2014: 51):

- There is a structure S such that S contains kind K as a position and K s perform function F in S .
- Members of K have instances of properties G_1, \dots, G_n .
- K s perform F though having instances of G_1, \dots, G_n .
- Anchors kind K such that----->
- For all persons P , the fact that P has instances of properties G_1, \dots, G_n grounds the fact that P is a member of human social kind K .

That the ontology of human social kinds offered here clearly reveals the anchors for their grounding conditions is taken to be a mark in favor of the ontology. Nevertheless, it is not argued that this is

the only possible way in which human social kinds might have their grounding conditions anchored, since different accounts of these social kinds may entail different anchoring schemas.

Rescuing Individualism, or How Supervenience May Survive the Materialist Challenge

Francesco Guala—Rescuing Individualism, or How Supervenience May Survive the Materialist Challenge
University of Milan

Contemporary defences of ontological individualism rely on supervenience and multiple realization. The classic reductionist requirement of a one-to-one correspondence (identity relation) between S-properties and I-properties has been replaced by the idea that each S-level property may be instantiated in a number of different ways at the I-level (e.g. Fodor 1974, Currie 1984, Sawyer 2002). This form of non-reductive individualism has been recently challenged by arguments that exploit the dependence of many social facts on material facts – i.e. facts that are not even about humans, let alone individuals (Epstein 2015, Haslanger 2017). One way to resist to this ‘materialist’ argument is to reply that the essence of individualism is the rejection of ‘spooky’ entities, properties, or facts. The claim that social facts supervene upon individualistic and material facts does not postulate the existence of anything ‘spooky’ and thus is not a counterexample to individualism understood this way (e.g. Hindriks 2013). Here I propose another strategy: granting the materialist her definition of individualism, I argue that individualists are committed only to the claim that some social properties supervene on individualistic ones – namely those properties that occur in bona fide social-level causal generalisations (call them ‘social kind properties’).

This point is often obscured by the contemporary shift toward the analysis of social ‘objects’ or ‘facts’, without further specification. Such a shift facilitates the anti-individualistic task of finding counterexamples, because there undoubtedly are some social properties that supervene on material facts. But it is by no means obvious that such properties are projectable or that they occur in social science generalizations, as genuine social kind properties should. For example: although G. W. Bush’s victory in the 2000 presidential elections depended crucially on the number of ‘hanging chads’ (Epstein 2009), the latter do not help explain any interesting social science generalisation that involves electoral competitions in democratic countries. If ‘being a hanging chad’ is irrelevant for social science, then it is not a counterexample for non-reductive ontological individualism.

Another way to put it is that individualism has always involved the consideration of two dimensions (think of Fodor’s famous diagram, or Coleman’s boat): (i) a ‘horizontal’ dimension depicting causal relations between properties at each level (social properties at the top, individual properties at the bottom); and (ii) a ‘vertical’ dimension mapping the metaphysical relation that holds between properties across the two levels. Like much contemporary work in social metaphysics, the materialist challenge focuses on the vertical dimension only, ignoring the horizontal one. But the horizontal dimension is crucial to identify the social properties that are supposed to supervene on individualistic ones. Ignoring the horizontal dimension makes the task of finding counterexamples too easy, and individualists should resist such a move.

In the last part of the paper I discuss and dismiss two possible materialist counter-attacks: I show that that my argument does not conflate ontological with methodological individualism, and that materialists are committed to a deeply revisionist picture of social science.

Political Responsibility Without Participation and Vice Versa: A Comment on Iris Marion Young's Social Connection Model

*Mattias Gunnemyr—Political Responsibility Without Participation and Vice Versa
Lund University*

One of Young's major aims in *Responsibility for Justice* (2011) is to make sense of a central claim of the anti-sweatshop movement: that agents such as consumers, retailers and global garment companies have a responsibility for the conditions under which apparel items are produced in many parts of the world. The result of her enquiry is the social connection model. This model involves the claim that all those and only those who participate in an unjust structural process, such as the global garment industry, have a forward-looking responsibility to ameliorate this process. Young often calls this political responsibility. The social connection model also states that the nature and degree of an agent's political responsibility vis-à-vis a certain structural process can be roughly determined by considering the power, privilege, interest and collective ability – i.e. the capacity (cf. Barry, 2005) – this agent has in virtue of participating in the process.

In this paper, I question the idea that we can determine who is responsible for redressing the unjust structural process *S* by considering who participates in *S*, while we can determine the nature and degree of each agent's political responsibility by considering entirely different factors. By way of giving examples, I show that the considerations concerning the capacity to redress injustice are relevant for deciding who has political responsibility. I conclude that participation in an unjust structural process is not a necessary condition for having political responsibility. I also consider some alternative interpretations of the social connection model, such as the ones suggested by Nussbaum (2011) and McKeown (2015), and argue that they cannot avoid this problem.

Furthermore, participation in an unjust structural process is insufficient for having political responsibility, at least if we assume what I take to be a natural understanding of forward-looking responsibility: that those who have a forward-looking responsibility for an unjust structural process ought to take steps in order to redress this process. That participation is insufficient for having political responsibility in that sense seems most obvious if we consider those who are most disadvantaged by unjust structural processes, like the workers producing clothes for the global garment industry. Even though some of them could join a union or broadcast their situation to others, there might be workers who participate in the process without having any opportunities of this kind. The risks of harassments, unemployment and ensuing poverty, for themselves and for their families, are simply too high (Gould, 2009, and McKeown, 2015, have discussed similar issues in relation to the social connection model).

Finally, I question the idea that the nature and degree of an agent's political responsibility can be decided only by considering the agent's capacity to redress injustice. I argue that an agent's nature and degree of participation might influence his or her nature and degree of political responsibility. The upshot is that considerations concerning both participation and capacity are relevant for deciding the initial question of who has political responsibility, and not only for determining its nature and degree.

A Taxonomy of Prospective and Retrospective Moral Responsibility in Collective Contexts

Niels de Haan—*A Taxonomy of Prospective and Retrospective Moral Responsibility in Collective Contexts*
University of Vienna

In this article I present a taxonomy of prospective and retrospective moral responsibility by looking at the various blameworthy actions and omissions in collective contexts. I argue that any unified theory of (collective) moral responsibility has trouble preserving three principles that seem uncontroversial for individual cases: (i) only moral agents can be subjects of moral imperatives; (ii) ought-implies-can; and (iii) retrospective moral responsibility presupposes prospective moral responsibility. I equate prospective moral responsibility with moral duties and moral obligations, and retrospective moral responsibility with moral blameworthiness. Thus, one can only be blamed for doing X, if one had a duty not to X (and vice versa).

First, I set out three types of retrospective responsibility. Individual responsibility: for an agent to be morally responsible the agent must be an apt candidate deserving of blame for that action or omission. Collective (or non-distributive) responsibility is the blameworthiness for an action or omission of a collective agent. I assume that, unlike unstructured collectives, structured collectives (states, corporations) qualify as collective moral agents (French 1984; List & Pettit, 2011). Shared (or distributed) responsibility refers to the aggregated responsibilities of the members of a collectivity, it need not presuppose the existence of a cohesive group (May 1992).

Next, I identify four types of blameworthy outcomes or actions: divisible aggregated outcomes, non-divisible aggregated outcomes, participatory collective actions, and joint actions. The first two are unintended, the latter intended. If a collective agent was involved in any of these types, the collective agent is collectively responsible, presupposing a collective duty. Without a collective agent present, I argue, the responsibility of all the agents involved is shared concerning the last three types, which presupposes the violation of individual duties, given that agents can refrain from participation. Thus, we need to loosen (iii), it need not only pertain to a strict type-type relation, as shared responsibility can presuppose the violation of individual duties.

I further identify four types of blameworthy inactions and omissions: incremental collective actions, threshold collective actions, responsive joint actions, and complex joint actions. Omissions of joint actions may pressurize (i) and (ii) if there was no (willing) collective agent present, given that agents are incapable of performing a joint action alone. However, for complex joint actions, the solution was not reasonably clear, requiring plan-like deliberation and role-distribution, therefore agents share responsibility for failing to organize (Held 1970), which presupposes collectivization duties (Collins 2013). These are individual duties to take responsive steps towards the formation of a collective agent, that incurs the collective duty. Yet with responsive joint actions the agents were already capable of performing the action together, the solution was obvious, requiring no organization, thus they share responsibility for the omission. Unless capacity for joint intentional action is equated with collective moral agency, (i) or (ii) must be revised. I argue that loosening (ii) is more attractive, accepting joint duties, where ought-implies-can pertains to a plurality of agents (Schwenkenbecher 2014). Finally, I anticipate and refute objections that we should loosen (i) rather than (ii).

Constitution of Social Kinds as Generic Identity

Jani Hakkarainen—*Constitution of Social Kinds As Generic Identity*

University of Tampere

S. Haslanger (1995; 2012) and R. Mallon (2008; 2016) distinguish crucially constitutive social construction from causal social construction. For example, Haslanger thinks that the social kind of women is constituted by having, to put it roughly, a subordinate and oppressive social position in relation to men (2012, 132–3).

However, neither Haslanger nor Mallon gives any detailed account of the relevant notion of constitution. I propose that it is to be considered identity when the latter is understood as the generic form of generalized identity.

The explicit notion of generalized identity is a newcomer in philosophy although its plausible examples are familiar: for instance, “for an entity to be a water molecule is for it to be an H₂O molecule”. A. Rayo (2013), Ø. Linnebo (2014), who coined the term, C. Dorr (2016), F. Correia (2017) and F. Correia & A. Skiles (2017) have done groundbreaking work on generalized identity.

I follow Correia and Skiles and consider generalized identity analogous to the familiar objectual identity (e.g. “Hesperus is Phosphorous”). Correia & Skiles (2017, 3) express generalized identity “with an operator, \equiv , indexed by zero or more variables, which takes two open or closed sentences and yields another.” Generic identity is generalized identity of the form “for an entity to be F is for it to be G” in the monadic case ($Fx \equiv x Gx$), which can be generalized into polyadic cases that involve relational predicates. Generic identity, just like objectual identity, is reflexive, symmetric and transitive.

My proposal is that the constitution of social kinds is generic identity. For example, for an entity to be a woman is for it to be (roughly) subordinately and oppressively socially positioned when this statement is considered expressing a generic identity. Its generalization gives us an elegant and simple understanding of the so far mysterious constitution of social kinds. For an entity to be a member of a social kind is for it to be such and such when “is” is construed as “the is of generic identity”.

The symmetry of generic identity does cause any problems for my proposal although constitution is typically considered asymmetric in general metaphysics. First, in the case of womanhood, \equiv is flanked by a conjunction on the right-hand side. There is no problem in saying that for an entity to be subordinately and oppressively socially positioned is for it to be a woman. Secondly, generic identity allows of representational differences between the left-hand side and the right-hand side, as well as in the objectual identity “Hesperus is Phosphorus”. So the right-hand side may be informative about the left-hand side and hence representationally more fundamental. Thus, the constitution of social kinds may very well be symmetric.

Furthermore, the right-hand side of a generic identity statement can bring with it a hierarchy of social kinds since it can involve asymmetric relational predicates. Generic identity does not have to be restricted to the monadic form. Women may very well be in a subordinate position to men in a way Haslanger maintains.

Team Reasoning and the Irreducibility of Collective Intentionality

Raul Hakli—Team Reasoning and the Irreducibility of Collective Intentionality
University of Helsinki

The paper defends an understanding of collective intentionality in terms of irreducible we-intentions. It presents an argument against reduction of we-intentions on the basis that there are cases in which such reductions necessarily lead to a loss in inferential power. It argues that team reasoning, when formulated in terms of we-intentions, in cases like Hi-Lo games allows the reasoning agents to reach conclusions that become unreachable if premises expressing we-intentions are replaced by corresponding premisses expressing individuals' intentions concerning group actions or individuals' intentions concerning their own actions.

A Platonic Perspective in Husserl's Social Ontology

Mary Elizabeth Halper—A Platonic Perspective in Husserl's Social Ontology
The Catholic University of America

Recently, some efforts have been made to give a more subtle account of Edmund Husserl's unique position on the spectrum of ontological individualism and ontological holism. This paper presumes the worth of these efforts and attempts to contribute to the development of Husserl's social ontology. More specifically, the aim of this paper is to bring the Platonic perspective found in Husserl's mereology to bear on Husserl's social ontology. In order to motivate this aim, the first part of the paper briefly presents the case that Husserl's account of parts and wholes, that is, his mereology, crucially informs his social ontology. For this case, I draw heavily from Petranovich 2017. Given that a better understanding Husserl's theory of parts and wholes will better help us situate his social ontology vis-à-vis individualism and holism, I turn to the task of deepening some aspects of Husserl's mereology.

The second and most involved part of the paper conducts an examination of the debt Husserl's mereology owes to Platonism. Drawing from Hopkins 2010, I emphasize the Platonic background to Husserl's theory of eidos and thereby develop what Hopkins calls Husserl's "arithmological theory" of eidetic wholes. On this theory, an eidos is a unity over a multiplicity that is responsible for the commonality of its parts without having anything in common with its parts. I then recast Husserl's mereology in terms of this arithmological theory, focusing especially on Husserl's distinction between parts that are "pieces" or parts that are "moments," and the relationships of "founding" and "mediation" within a whole. In this way I use the arithmological account to theorize a distinct way of being a whole.

The third part of the paper applies this distinct way of being a whole to Husserl's social ontology. Here I suggest some ways in which this sense of being a whole can address questions that arise about Husserl's conception that is the whole of (some) social communities. In particular I focus on how the sense of being a whole developed here may affect how we understand the intersubjectivity and Umwelt ("surrounding world") that, in some sense, belongs to the social community. After developing this account I apply it to some familiar puzzles that arise for individualism and holism, and thereby further clarify the unique character of Husserl's social ontology.

Structuralism and Individualism Functioning Together

Jasper Heaton—Structuralism and Individualism Functioning Together
The University of British Columbia

Structuralists criticize social explanations that account for oppression ultimately in terms of people's choices, psychological states, and actions. Structuralism arose as theorists began paying closer attention to the ways social institutions shape people's lives. Structuralists consider how institutions limit the spaces of possible options that people have and impact the architecture of their choices. In this way, institutions constrain and enable people, channeling them down gendered paths of life. For instance, company policies about the availability of maternity pay can impact the architecture of the option-spaces people choose from when thinking about childcare, drawing more women to take on the role. Structuralists thus identified a structuring function that social institutions perform.

At times Haslanger and other structuralists veer towards identifying social institutions with the structuring function they perform. I argue that we must avoid identifying the structuring function with the institutional ontology performing that function, because that function is sometimes performed by individualistic factors. If we consider structuring as a function, we can capture the important structuring roles that individualistic factors sometimes perform.

Though her focus is not so limited, many of Haslanger's examples explore structuring in industrialized contexts, with states and corporations. In other contexts different sorts of social units are salient. Marriages, and friendships, for instance, are potentially oppressive contexts, and in ways that involve individualistic factors impacting the architecture of choice in actual situations. The gendered expectations of a romantic partner, coupled with forces that pressure her into staying in a relationship, structure the possible choices available to a woman. Masculine emotional repression severely stunts the range of emotional options many men have; this is a real, psychological lack that plays a structuring role for the emotional options men take in their relationships. Individualistic factors of course do not exhaust the structural causes of such situations; they stand among them.

Once we acknowledge structuring as a function, we see that institutional factors are themselves germane to structuring explanations, and that individualistic factors play a role. This, ultimately, is to acknowledge that structuring factors are themselves apt for their own structural explanations. The architecture of choice for the proposal of legislation is such an example. For instance, we should consider structuring explanations for why bathroom bills, rather than weapons bans, are proposed for the good of the public's safety. Those explanation will have, among their explanans, the preferences of voters.

My point is not to recognize a class of oppressive situations structured solely by individualistic factors, but to come closer to capturing the complexity of the ways social situations can be oppressive. Perhaps some situations are structured more by individualistic factors, others by institutional factors, but most actual situations will be structured by a confluence of factors, institutional, individualistic, and more. Further, this makes clear that institutional factors far from exhaust the the constituents of the social milieu that structure our lives. Social milieu decompose

into a range of factors, all of which are, in different contexts, apt for structuring explanation or structuring role.

Writing Together: What the President's Daily Brief Can Reveal About the Significance of Collective Written Testimony

Marcus Hedahl—*Writing Together: What the President's Daily Brief Can Reveal About the Significance of Collective Written Testimony*
U.S. Naval Academy

In February of 2018, The Washington Post reported that President Donald Trump is the first President in fifty years who has chosen not to read the President's Daily Brief, the distillation of information put together by U.S. intelligence agencies (Leoning, Harris, and Haffe 2018). While the President has been roundly disparaged for his lack of analytical rigor and intellectual curiosity (See, for example, Graham 2018), he can also be criticized for turning away from an under-analyzed, but important source of knowledge: collective written testimony. In this paper, I look to collective epistemology to analyze the ways in which written collective testimony differs from verbal collective testimony, in order to consider how that distinction can shed light both on the feasibility of deflationary accounts of collective testimony and on the significance of the vulnerability problem regarding collective testimony.

In Sections I and II, I briefly explore why the predominant paradigm within the collective testimony literature has remained verbal testimony an agent makes on behalf of a collective, a fact that can be explained in part because of the adaptation from individual testimony to the collective case, in part because doing so less directly confronts some of the problematic issues of collective epistemology (e.g. whether there must be a corresponding collective belief (Cohen 1989)), and in part because testimony requires elements that are more often found in verbal communication than in written communication (e.g. Testimony may have to be interpersonal in a way writing typically is not, testimony may require uptake that is not necessarily guaranteed when writing (Fricker 2012)). In Sections III and IV, while trying to remain agnostic about an ultimate theory of testimony, I demonstrate how written testimony can, in very specific circumstances, meet the criteria required for testimony. In Section V, I argue that in Bratman's (2014), Gilbert's (2013), or Tumela's (2016) frameworks for analysing collective intentional action, collectively written testimony is more likely to be the output of genuinely collective intentional action than collective oral testimony would be. In Section VI, I build on this insight, analyzing two reasons why considering collective written testimony as a distinctive form of testimony is theoretically significant. First, since collective written testimony is more likely to be the output of genuine collective intentional action, adding writing as another paradigmatic form of testimony demonstrates a serious shortcoming for deflationary accounts of collective testimony (e.g. Lackey 2014). Second, since collective testimony can be the product of genuine collective intentional action, in certain conditions, it can become less susceptible to the vulnerability problem than individual testimony would be in those same circumstances (Faulkner 2002; Alder 2002). Collective written testimony thereby offers a distinct form of epistemic access that cannot be reduced to individual testimony, the testimony of a supermajority (c.f. List 2014), or to collective oral testimony an agent makes on behalf of a collective (c.f. Lackey 2014). The hope is that proceeding in this manner might offer a novel way to move past what is sometimes characterized as a stalemate in the debate between summativists and non-summativists about the possibility of collective belief (See, for example, Gilbert & Pilchman 2014: 210).

Throughout the paper, I use two examples to highlight the theoretical points under consideration: The President's daily intelligence briefing, and the National Reconnaissance Office's report on the future of satellite ground systems, the latter of which had thousands of individual agents and dozens of collective actors acting as contributors to the final product (to include the author of this paper). These examples help highlight not only the advantages but also the limitations of this distinct form of epistemic access.

Derogatives: Beyond Insults and Slurs

Cassie Herbert—*Derogatives: Beyond Insults and Slurs*
Hobart & William Smith Colleges

There is a great deal of variety amongst harmful words, both in terms of their history, the sorts of relations they call forth, the contexts in which they function powerfully, and the way they're used and deployed in everyday speech. The terms 'slut,' 'basic,' and 'Becky,' for example, all target women on the basis of their gender and so (on some readings) might all be simply classified as slurs. Yet none of these are 'typical' slurs - for example, none of them are wholly taboo. At the same time, they all have distinctly different meanings and effects, and are used in quite different ways. 'Slut' roughly means a woman who is promiscuous. 'Basic' describes a young, white woman whose interests and personality are uninteresting. And 'Becky' marks out white women who profit at the expense of people of color. All of these terms are both used on the basis of gender and have at least the veneer of being warranted in virtue of a woman's behaviors and activities - whether that's sexual activity, a certain set of derided interests, or a kind of social climbing that comes at the expense of people of color. Our existing conceptual framework for thinking about slurs, however, simply can't capture these sorts of differences, and even expanding the framework to include both slurs and insults doesn't offer much help.

I argue that our category of 'slurs' is both too narrow and too coarse. Too much gets left out of this category, and at the same time, the category isn't fine-grained enough to pick out the differences between harmful terms. Instead of focusing on whether or not something is a slur, I argue that we should start thinking in terms of 'derogatives.' The category of derogatives has a number of continuums along which a term may be placed. Derogative terms encompass paradigmatic slurs, which often fall at the extreme ends of these continuums. But derogatives also include terms with more complex locations, terms that essentialize their targets without being taboo, or terms that don't elicit such deeply entrenched systems of oppression. Derogatives call for attending to the intersectional aspects of terms, the ways some terms target multiple aspects of identity or do different work when directed at differently positioned members of a group. Just how a derogative is located on this multidimensional continuum impacts how it functions in the world. Acceptable derogatives are positioned to be used more frequently than taboo derogatives, and this gives them a particularly insidious power. Community-specific derogatives may be incomprehensible to those outside the community, and so have more constrained scope of use, but may function especially powerfully within the communities in which they're deployed specifically because of their status as insider-speech. Finally, I argue that while slurs are always wrong to use, derogatives used by marginalized communities to resist systems of oppression may at times be warranted. This expanded taxonomy of harmful terms will equip us to better understand the rich variations amongst slurs and other derogatives, and the distinct harms such terms can enact.

Re-Bunking Corporate Agency

Kendy Hess—Re-Bunking Corporate Agency
College of the Holy Cross

In this paper I suggest that the corporate agency literature in general – and the primary first generation accounts in particular – have not paid enough attention to the metaphysics of the entities in question. The first generation accounts from French, Werhane, and List and Pettit (like the more collectivist accounts from Gilbert, Miller, and Bratman often included in the literature) have focused on corporate agency to the exclusion of corporate agents. Each proponent has presented a single mechanism by which a collective could exercise agency, and the general silence regarding other properties of the collective has left the impression that this (single) mechanism defines the entity – that this single mechanism is the sole source of its unity, cohesion, endurance, intentionality, action, and identity. Critics have read the accounts in this way and responded with an entirely appropriate skepticism. Velasquez, Wall, Narveson, Miller and Mäkelä Rönnegard, Ludwig and others have charged that such an entity is “ghostly” and “mysterious,” an “incorporeal” “hovering entity” too flimsy to support the robust capacities proponents claim for it. To those criticisms I add my own: that this presentation does a disservice to the holist position more generally, failing to acknowledge the deep integration of the corporate agent, the undeniable presence of multiple mechanisms for corporate agency, and the significant contributions that lower-level members make to both the content and the enactment of the corporate agent’s intentionality. The latter point especially has significant practical and normative implications.

The paper begins with a quick sketch of my own, metaphysically robust account of corporate agents. On my account, a corporate agent is unified and defined by a “structure” (Haslanger 2016) that embodies a rational point of view (Rovane 1998). This structure is a heterogeneous mix of components that typically includes (1) all of the formal, intentional mechanisms identified by the first generation proponents listed above, (2) additional, crucial mechanisms that allow for broadly distributed and unintended decision-making, and (3) the tacit, informal mechanisms that play such a significant role in social life. Recognizing the richness and power of this immanent structure allows us to recognize the corporate agent as a robust whole – something that qualifies as a material object on standard accounts (Baker 2002, van Inwagen 1995), that possesses agency without being entirely defined or determined by it. Among other things, it allows us to demonstrate that

- (1) the corporate agent necessarily possesses a material aspect (nothing ghostly or mysterious about it);
- (2) the distributed material aspect of the corporate agent is familiar from other “scattered objects” (Cartwright 1975) and not problematic; and
- (3) the reliance on human intentionality for “fastenation” (Markosian 1998) is unique to social entities but not problematic for claims of existence or agency.

On this account, the corporate agent is a material entity much more similar to the paradigm human agent than is typically recognized by either proponents or critics. This approach allows us to respond to all of the criticisms listed above, and more besides.

The Duty to Join Forces

Frank Hindriks—*The Duty to Join Forces*

University of Groningen

The hallmark of collective responsibility is the possibility that a collective has some responsibility even though none of its members bear a correlative responsibility. This ‘Irreducibility Thesis’, as I have called it (Hindriks 2009), has been defended not only for agential collectives, but also for collectives that are not agents. For instance, Joel Feinberg (1968) argues that the passengers involved in the Jesse James train robbery had a collective duty to rise up as one and subdue Jesse James, even though none of the individual passengers was at fault. In this paper, I do two things. (1) I argue that the Irreducibility Thesis does not apply to non-agential collectives, because such collectives do not have irreducible responsibilities. (2) I propose that their members have what I call ‘a duty to join forces’ (cf. Held 1970 and Collins 2013 for a similar proposal for agential collectives). This is an obligation to form a joint intention and bring about the relevant outcome together.

(1) The Irreducibility Thesis for non-agential collectives has been defended by invoking non-existent hero’s, agents that have yet to be formed, and interrelated capacities that exist in a morally ideal world (Feinberg 1968; Wringe 2010; Björnsson 2014). The core of my criticism of the Irreducibility Thesis concerns the fact that, by definition, non-agential collectives are not agents. As only agents can act for reasons, only agents can have obligations. This in turn entails that only agents can have responsibilities. Because of this, it makes no sense to attribute responsibilities to non-agential collectives. I argue that this line of argument generalizes to the three positions mentioned that try to get around this.

(2) I go on to argue that individual members of non-agential collectives can have a duty to join forces with the others. This is the duty of an individual to take steps towards forming a joint intention. It requires of an individual to take initiative and approach someone else to form such an intention, or to follow suit when someone else approaches her. The duty to join forces is owned by individual agents. In this respect, it is an individual responsibility. At the same time, it is collective in two respects. First, an individual agent has this duty only if some other agent has it as well. Second, its content is collective in that it concerns the formation of a joint intention.

Only once those bounded by it successfully form a joint intention do the members incur the obligation to act on it – to perform the requisite joint action. This entails that, if the relevant individuals do not form a joint intention to help, the duty to do so never materializes. Anne Schwenkenbecher’s (2014) objects to this claiming that this entails that certain responsibilities remain ‘unallocated’. I argue that we should bite the bullet and accept it, as the alternative of attributing responsibilities to non-agential collectives is even more unpalatable.

Socially Constructed Kinds

Frank Hindriks—*Socially Constructed Kinds*
University of Groningen

It has often been argued that social objects and properties can be natural or real kinds. This is what I call ‘the real kind thesis’. The main challenge to this thesis is that social entities are in some sense mind-dependent. This is not an unsurmountable obstacle as long as the relevant entities have causal properties. In this paper, I use the rules-in-equilibrium account of social structures to provide an improved defense of the real kind thesis. The rules-in-equilibrium account serves (1) to make precise the way in which social entities are mind-dependent, and (2) to explain why social entities have causal properties. The argument makes use of recent philosophical theories of social construction as well as of core scientific insights into the explanation of social behavior.

(1) Mind-dependence. I argue that the vast majority of social entities are acceptance-dependent. This means that certain people accept particular entities as having some social status – rivers as borders, pieces of paper as money, and people in particular uniforms as police officers. Such statuses consist of rules that feature activities that are characteristic of the relevant entities – don’t cross the river, use the pieces of paper for buying goods, and listen to people wearing those uniforms. Because of this, people tend to respond to certain triggering conditions involving those entities by engaging in characteristic activities. Hence, many social entities are status entities that depend on actions as well as attitudes, and that are socially constructed in both the causal and constitutive sense.

(2) Causality. Using insights from social scientific theories of conventions and institutions, I argue how acceptance-dependence can in fact explain why social entities have causal properties. The most widely accepted explanation is that people often have an incentive to participate in existing behavioral regularities simply because others do so. More technically, they have a preference for engaging in activities that are characteristic of a status that is conditional on others doing so. This means that strategies to conform constitute equilibria. In this way, status entities cause behavioral regularities that are self-reinforcing. Because of this, those social entities into stable and causally relevant status entities.

I use this framework to address and resolve recent controversies about abstract social kinds (Smith, Searle), conventional social kinds (Khalidi), and infallibilism about social kinds (Thomasson).

Multi-Dimensional Personhood and Attributing Responsibility to Group Agents

Onni Hirvonen—Multi-Dimensional Personhood and Attributing Responsibility to Group Agents

University of Jyväskylä

It is commonplace to find groups in such roles in our everyday language on politics, social sciences, and law that requires attributing beliefs, agency, and moral personhood to them. At the same time it is also common to hold onto a disanalogy between (natural) human persons and (artificial) group persons (see Smith 2018). Though it is partly unclear what fully-fledged personhood entails, we do not, for example, easily grant voting rights or rights to social security for group agents.

This paper argues that a fruitful way of making sense of the disanalogy is to understand the concept of personhood as a concept with multiple dimensions. Firstly, personhood includes both psychological and social conditions (see Laitinen 2007). This follows the common intuition that, for example, attribution of moral personhood requires that one is a capable member of a moral community and, thus, at least capable of moral reasoning. This account challenges simple performative views of personhood (e.g. List and Pettit 2011) by accepting that there are certain psychological or ‘intrinsic’ features like rationality, capability to reason, and having linguistic skills that are necessary for being recognized as a person.

Secondly, personhood can be understood as a historical-political concept. That means that the personhood-constituting relationships (or performances) have developed historically and that they may also come in different forms. This is an idea that has been extensively argued for by the Hegelian recognition-theorist who often also accept that different forms of person-constituting social recognition have different psychological/intrinsic conditions. This paper aims to combine these two ideas to introduce a view in which personhood has historically changing multiple aspects, and in which separate psychological capabilities and social recognition intertwine.

The multi-dimensional view of personhood can be argued to be problematic because it does not seem to offer clear criteria to distinguish persons from non-persons. This paper claims that this worry need not be damning as there might just be no strict metaphysical delineation between persons and non-persons available. Questions like “who counts as a person” and “what are the relevant conditions that need to be fulfilled to get recognition as a relevant type of person” may well be partly open-ended and the answers we give to them can change along with our social practices. However, adopting the multi-dimensional view of personhood allows some of the currently contested cases – like social robots and group agents – to be understood as at least partial persons (if not full persons in the exact same sense as most adult human agents are) all the while holding onto the intuitions that (a) responsibility can only be attributed to persons, (b) personhood requires certain agential capabilities, and (c) there are relevant differences in the capabilities of human and group agents.

Non-Agential Collectives and Consumer Complicity for Climate Change

Säde Hormio—Non-Agential Collectives and Consumer Complicity for Climate Change
University of Helsinki

Arguably no state is currently doing as much as they should if we are to keep the global temperature rise limited to 1.5 degrees Celsius. As a fair collective scheme to deal with climate change is currently lacking, what should individuals do in the face of a collective failure to mitigate climate change? The vast majority of our emissions are linked to a collective context of one kind or another, and to change these patterns of emissions we will need to change not just structures, but also certain collective patterns of behaviour, for example the way we act as consumers.

Non-agential collectives are a set of individuals who lack an established group structure, but are nevertheless connected by some common feature, whether that feature is persisting or temporary. Examples include different kinds of interest groups and social movements, sets of people who happen to be at a certain place at a certain time, but also groups of people that together cause harm in aggregate and often unintendedly, like consumers. What makes these cases collective is the context. While the individuals might act independently, they do not act in isolation. That is to say, the individual acts take part in a collective context and their meaning and significance can only be fully captured by analysing the collective level.

I allow that we can hold constituents of non-agential collectives blameworthy in the backward-looking sense under certain conditions. However, I deny — contrary to Isaacs (2011) and Wringe (2010) — that you can place obligations, putative or actual, on non-agential collectives. Instead I argue that our duties qua non-agential collectives are always interdependent individual duties. Instead of obligations of putative collective agents, I suggest that we can have shared responsibility as constituents of non-agential collectives. Collective responsibility does not need to be based on collective agency: shared intentions (Kutz 2000, Isaacs 2011, May 1992) are an alternative basis for collective responsibility.

Nonetheless, I grant that it can be useful to discuss non-agential collectives in some cases, as it can help us to make better sense of our complicity for collectively caused harms. More specifically, it can make us appreciate the different structures and systems that we are part of, and how we are complicit in upholding and recreating these. When it comes to climate change, instead of participatory intentions, the basis of complicity is quasi-participatory (Kutz 2000). Thus the complicity of consumers is not based on some intention to emit, or intention to participate in climate change harms, etc. (cf. Lawford-Smith 2017). Instead, it is based on the quasi-participatory intention to partake in a consumer lifestyle, a lifestyle that has excessive emissions as an unintended, but foreseeable consequence.

Individualism, Psychology, and Collective Intentionality

Graham Hubbs—Individualism, Psychology, and Collective Intentionality
University of Idaho

Accounts of collective intentionality tend to be individualistic and psychologistic. The commonplace strategy is to start from the presumption that all intentional phenomena are to be explained in terms of the psychological states of individuals and then to characterize collective intentionality in terms of the propositional contents of a collection of individual psychological states. The goal of the present paper is to challenge this approach. The argument is primarily critical, but it belongs to a broader project of explaining collective intentionality in terms that are neither individualistic or psychologistic.

I begin by outlining Michael Bratman's account of collective intentionality, explaining why it can serve as an exemplar of the individualistic psychologistic strategy. I then review a challenge that Bratman has wrestled with over the years, viz., how on his view one person can, without being in a position of authority, "settle" the intentions of another. Bratman needs to meet this challenge for the idea of an individual psychological state with a collective action as its propositional content to be intelligible. He addresses this challenge by claiming that a person settles another's intention if the first forms the intention that the second do something and can predict with a high degree of accuracy that the second's knowledge of the first's intention causes the second to intend to do what the first intends. This means that on Bratman's view settling a collective intention is always a contingent matter; it is always only a matter of contingency that the second person doesn't rebel, so the first can only ever contingently settle the second's intention.

I argue that when we combine this line of thought with reflection on the instrumental and embedded structure of intentional action, we see a basic problem with Bratman's conception of settling, a problem that threatens any individualistic psychologistic view. Many (if not most) intentional actions comprise instrumental steps that are themselves intentional actions. Sometimes, in the course of acting, an unforeseen step arises that the agent or agents find challenging and that causes them to question whether to continue trying to execute their broader intention. Even if Bratman's view can make sense of one person initially settling the intentions of another, it cannot articulate what is required for a collective to persist with or to quit that intention in the face of an unexpected challenging step. The reason: as noted above, on Bratman's view one can settle another's intention only contingently, but the connection between an act of decision and the maintenance or abandonment of an intention is non-contingent. This connection is, in a sense I adopt and explicate from Richard Moran, immediate. I argue that neither Bratman nor anyone else can account for this immediacy while maintaining the individualistic psychologistic strategy of explaining collective intentionality, and on that basis, I conclude that we should give up the strategy.

The Social Ontology of African Communalism and Personhood

Polycarp Ikuenobe—The Social Ontology of African Communalism and Personhood
Kent State University

I examine the social ontology of African communities, in order to show how the communal traditional view of the social world, reality, and mode of existence, provides the foundation and grounding for a conception of personhood. African communalism, on the one hand, indicates a social ontology of personhood, which requires the social process of internalizing communal norms and beliefs. On the other hand, it involves individuals' ability to use communal norms to reason and guide one's conduct, in order to create the social ontology of harmonious communal living. The communal norms that provide the grounding for one's personhood exemplify collective destiny, identity, solidarity, values, beliefs, intentionality, and reasoning, which provide the foundation for collective action. Thus, one's individual intentionality and actions are a part of, and contingent on, communal intentionality and actions.

I examine John Mbiti's famous rendition of the social ontology of African communal view of personhood that is captured by the cryptic statement: "I am because we are, and because we are therefore I am." I analyze this statement as an account of African social facts and processes. It indicates communal beliefs, functions, and norms with deontic powers that shape individuals' selfhood as social-moral objects, their beliefs, functions, status, obligations, reasoning, and actions. I argue that this statement also involves how individuals use rationality to create and justify African communal social reality, functions, practices, constitutive rules and their deontic power. Thus, communalism indicates Africans' collective acceptance of beliefs and constitutive rules that the other-regarding value of harmonious communal living is good, and that the good of communal living is for the good of individual flourishing and well-being. These constitutive rules and their collective acceptance and intentionality provide the basis for individual values and reasons for actions that are directed at communal harmony and solidarity.

I examine Ifeanyi Menkiti's distinction between a minimal conception and a maximal conception of personhood. The African idea of communalism indicates that a maximal conception personhood, as a social object, is grounded in the idea of communal harmony, solidarity, and mutuality. I argue that a proper understanding of his maximal conception of personhood must involve an understanding of the social processes of African communalism.

This communal idea is also expressed by the notion of Ubuntu, which indicates the idea that, a person is a person through other persons. This idea is usually understood as a communal norm or value that, a person's ultimate aim is to act to improve one's own well-being, and it is best achieved by communal harmonious living and solidarity with other persons. Thus, African communalism involves the social ontology of mutual dependence among individuals, and between individual and the community. I examine briefly the implication of the social ontology of African communalism and personhood for an African view of human rights, dignity, and autonomy.

Awareness, Dependency, and the Heart of Blame

Bernard Jackson—Awareness, Dependency, and the Heart of Blame
Regis College

Various arguments for reparations for those United States citizens of African descent whose ancestors were the victims of slavery have been forwarded for a number of years. The present essay holds that such arguments are dependent upon what has come to be called a “responsibility gap”: namely, a situation in which an unstructured collective or a collective agent is responsible without any of its individual members bearing a correlative responsibility. The reparations movement is motivated by the view that the effects of slavery still exist, and that the lives of the ancestors of the victims of slavery have been, to say the least, adversely affected. However, it is difficult to identify the unstructured collective or collective agent that has perpetuated this state of affairs. Is it white Americans? Is it the US government institutions? Or is it the Africans who colluded with the slave-traders centuries ago? But who specifically is to blame for the situation in which many African-Americans find themselves? Who are the members of this unstructured collective? Do it include all white Americans, including children? Are all US institutions guilty? Is there are a group of present-day Africans who benefit from actions perpetuated over three hundred years ago? As I see it, a responsibility gap implies that no one is to blame. After all, the bearer of blame is an unstructured collective. The present essay argues that a notion of responsibility in which no one is to blame is an empty notion. Thus, the concept of a responsibility gap should be discarded. I would like my submission to be included in the symposium on collective responsibility.

Arguing for the Possibility of Group Self-Deception

*Genevieve Jacot-Guillarmod—Arguing for the Possibility of Group Self-Deception
Rhodes University*

In group settings it is not always obvious who is ultimately to blame for unethical practices. In many legal cases involving corporations, courts have concluded that no individuals within the company can be held truly accountable, and so no one at all is held responsible (i.e. the notorious case of the Herald of Free Enterprise ferry [Pettit]).

There are many examples of groups who espouse certain views, or perform certain acts, that have, to varying degrees, a negative impact on their societies. Growing numbers of parents are choosing not to have their children vaccinated due, in part, to the claims and actions of anti-vaccination groups. Many corporations go to great lengths to present themselves as being environmentally friendly when in fact they are not, misleading the public. For governments across the globe corruption is endemic. One possible explanation for such situations is self-deception.

Recently it has been suggested that self-deception is something that affects groups as well as individuals (Trivers, Deweese-Boyd), however very little literature regarding the phenomenon exists. Given the prevalence of the view that any group phenomenon can be fully understood in terms of properties of a group's individual members, one may wonder whether or not group self-deception is genuinely distinguishable from individual self-deception. Yet given that certain groups wield a great deal of political, social and economic power, if groups are capable of self-deception there is room for things to go awry on a very large scale with potentially dire consequences.

I would like to suggest that (i) there are a number of situations in which it is useful (or even indispensable) to utilise talk of group-minds and (ii) that some of these situations also meet plausible sufficient conditions for attributing group self-deception.

One of the most obvious ways of determining whether or not groups can be self-deceived is to see whether or not groups are capable of meeting Alfred Mele's four minimally sufficient conditions for self-deception. I shall argue that they can.

Over and above this, with reference to Philip Pettit and Christian List's examination of 'discursive dilemmas', I also hope to draw attention to the fact that groups seem to offer a fertile setting for a very literal take on self-deception (something that causes trouble for an account of individual self-deception), offering an additional reason to think that it would be a mistake to reduce instances of group self-deception to talk of individual self-deception.

Finally, acknowledging that groups are capable of self-deception could have implications for how we think about collective responsibility. A closer examination of how self-deception is instantiated within groups could elucidate to what degree groups should be held accountable for actions that fall under its sway. Even if one wishes to argue that one is not morally responsible for self-deception and the acts that stem from it, it may still be useful to establish what the mechanics of group self-deception are in order to limit certain problematic behaviour that may arise as a result.

Power as a Social Ability—Towards a Philosophical Explication

Daniel James—Power as a Social Ability—Towards a Philosophical Explication
University of Duisburg-Essen

Power is both one of the most central and one of the most contested concepts in social and political theory. One long-standing debate in both fields concerns the nature of power: What is it for an agent to have power? According to an answer to this question more or less implicitly assumed in many theories, power is a particular kind of ability. The most extensive account of power in terms of abilities can be found in Philip Moriss's seminal book *Power: A Philosophical Analysis* (1987). However, insofar as Moriss's account relies on a version of what is commonly referred to as the conditional analysis of abilities, it inherits some of the problems of such analyses: Notoriously, they struggle to account for masked abilities and have trouble capturing the gradability as well as the context-sensitivity of ability ascriptions. These issues make such an analysis of abilities unsuitable for an account of power: For power in the social world is often masked and ascriptions of power are also gradable and highly context-sensitive. Apart from inheriting these problems from the conditional analysis of abilities, Moriss's account of power in terms of ability fails to adequately explain what kind of ability we ought to take power to be, and, hence, does not allow for distinguishing between abilities in general and powers in particular – a failure which, in turn, raises the question why we require a distinct concept of power in the first place.

The aim of this paper is to amend these shortcomings by explicating a distinct concept of power against the background of a more refined account of abilities. To achieve this aim, it will do three things: Firstly, it will provide a framework for abilities which accounts for gradability and masks and is therefore suited for the analysis of power. Crucial to such an account is to move away from an analysis in terms of a simple counterfactual and towards a view that is spelt out in terms of proportions of worlds (see Manley and Wasserman 2008). Secondly, it will proceed from the insight that agents typically possess certain abilities in view of certain facts (see Kratzer 1989, 1991). Against this background, it will lastly revisit the question what kind of ability we ought to take power to be and restate this question as follows: Which facts matter when it comes to an agent's power, rather than their abilities in general? The answer to this question will amount to an explication of the concept of power which is guided by the idea that this concept ought to be fruitful for social-scientific explanations: On this account, the facts which are explanatorily relevant to an agent's power are typically social in nature. Thus, power – understood as a distinct concept of particular interest to the social sciences – ought to be understood as a social ability an agent possesses only in virtue of certain social facts.

Legal Positivism and Varieties of Dependence

Bartosz Janik—*Legal Positivism and Varieties of Dependence*
Jagiellonian University

The basic tenet of the contemporary legal positivism is well expressed by S. Shapiro when he says:

(SPP) Legal facts are ultimately determined by social facts alone (Shapiro, 2013).

The problem is, of course, what the ‘determination’ in question consists in. Should one understand (SPP) as saying that

(S-SPP) Legal facts supervene on social facts,

or rather as:

(G-SPP) Legal facts are grounded by social facts?

Before this problem can be investigated, however, one preliminary question needs to be asked: what are legal and social facts? It is usually assumed that social facts are a matter of description, while legal facts are understood as pertaining to rights and duties of individuals, i.e. are normative. However, it is not fully certain whether such a clear-cut distinction between what is descriptive and what is normative is viable. For example, social facts are expressed in language, and language – at least according to the prevailing view – is a normative phenomenon. Therefore, although for the sake of the project it will be assumed that legal facts are normative, and social facts are descriptive, the complication indicated above will be taken into account. It is also possible that the utilization of the concepts of supervenience and grounding will shed some light on this controversy.

The major difference between supervenience and grounding is that the former may be formally defined (for example by saying that if two possible worlds are indistinguishable in relation to the subvenient property S, they are also indistinguishable in relation to the supervenient property P). In this context, the clause S-SPP may be rendered as “For any two possible worlds w_1 and w_2 if they are identical in terms of social facts, then they are also identical in terms of legal facts.” The relation of grounding, on the other hand, is treated as a primitive notion (i.e., one which cannot be defined formally, save in a partial and contextual way) and expresses metaphysical causation. It follows – at least according to some scholars – that supervenience, in contrast to grounding, has no explanatory power. As such, supervenience may turn out to be too weak a relation to be a criterion of choosing between non-positivist and positivist theories of law. Both non-positivists and positivists may accept that the law is a supervenient entity, while upholding their respective views pertaining to the nature of law. In other words, it seems *prima facie* likely that the relation of supervenience must be amended with additional formal elements in order to serve as a useful tool in constructing a legal ontology.

Grounding, on the other hand, is a stronger relation, and involves an explanatory component. From this perspective, it seems better suited to account for the relationship between legal and social facts. However, its weakness is linked with the primitive character of the discussed relation. Even if its explicitly stated formal features are intuitively plausible, as long as the relation in question is not fully formally defined, the conception of law it provides will be necessarily vague. Therefore, it seems that in its current, quasi-formal version grounding cannot constitute a fully adequate way of describing the relationship between legal and social facts.

The research hypothesis of the paper is that neither the relation of supervenience, nor that of

grounding provide a sufficient explanation of the kind of correspondence between legal and social facts (qua natural facts) as described in the contemporary legal positivism. Without further modifications they fail as the formal account of positivistic legal ontology.

The Formal and the Informal. Towards an Integrative Theory of the Social

*Ludger Jansen—The Formal and the Informal. Towards an Integrative Theory of the Social
RUB Bochum & University of Rostock, Germany*

Researchers in the social sciences have found it useful to distinguish between formal and informal groups or institutions. In contrast to informal groups, formal groups are created explicitly with clear structural roles which are explicitly conferred upon their bearers and govern relevant behavior within the group.

Theories in social ontology draw in different ways on this formal–informal dichotomy. For instance, Margaret Gilbert in her theory of plural subjects starts from everyday group phenomena, like two friends going for a walk, but she models these phenomena as establishing an explicit deontic framework. Two friends that are jointly committed for a walk are, in her theory, a group that is formal in a minimal way. John Searle, in contrast, starts from examples like presidents and money which involve explicit declarative acts. But in his theoretical framework it is rather implicit mental acts that are key, namely sufficiently many matching we-intentions within a group.

Clearly, (1) the informal is more basic than the formal, as we need language, a paradigm informal institution, to institute formal institutions. (2) Informal institutions are important to formal institutions, like bureaucratic hierarchies. (3) Often it is hard to distinguish between formal and informal institutions, for instance between a marriage (a paradigm of a formal institution) and the accompanying partnership (a paradigm of informal institutions, like friendship), though we know that there are both marriages without partnerships (e.g., in sham weddings or shortly before a divorce), and partnerships without marriages.

Nevertheless, I argue that the formal and the informal are wholly distinct categories: Formal institutions are created; they need an explicit datable act of establishment by which rights and obligations are assigned that constitute the institution. In contrast, informal institutions do not have an explicit datable act of establishment, but they come into being gradually. They simply develop and are not constituted by deontic entities, but by dispositions to act in certain regular ways.

Now the following objection may be raised against my categorical distinction between the formal and informal: As we have duties towards our partners and friends, deontic entities are not restricted to explicitly created institutions; hence, there is no need to draw a sharp distinguishing line between the formal and the informal. This objection fails because the duties of partnership have a quite different origin than the duties resulting from explicitly established institutions. Duties of marriage are socially established duties, like any that come about through contracts and promises. I argue that duties of friendship and partnership are derivative moral duties, either (1) as special cases of duties to other humans in general, or (2) due to the fact that partners have specific opportunities for actions which will be of specific significance for their partners.

As I will argue, monopolizing one side of the formal–informal dichotomy in a theory of the social world leads to a misrepresentation of the other side. I will indicate how an integrative theory of social ontology can do justice to the particularities of both formal and informal social entities.

Understanding Oppressive Ontologies

Katharine Jenkins—Understanding Oppressive Ontologies
University of Nottingham

This paper draws on recent work by Brian Epstein (2015) and Sally Haslanger (2016, 2017), among others, to offer an account of social ontology focused on two basic notions: social practices, and interest-relative explanatory kinds. It then demonstrates how this account (1) provides important resources for understanding the social ontologies involved in oppression, and (2) alerts us to a potential hazard of this kind of investigation.

Social practices involve picking out, or ‘framing’, certain kinds for differential treatment. These ‘conventional kinds’ need have no prior explanatory value, but they become explanatory once they are framed by social practices. Conventional kinds in turn function to organise and direct social practices, sometimes ranging far beyond the social practices which originally framed them. This process of ‘scaffolding’ gives rise to further explanatory kinds which exhibit causal regularities similar to those exhibited by the explanatory kinds of the natural sciences.

To illustrate, consider a situation where the only legal status a worker can have is either ‘employed’ or ‘self-employed’, but a gig-economy develops in which workers have working patterns typical of employed workers but are legally classed as self-employed and so have far fewer rights and worse pay. There are patterns of similarity between gig economy workers that are causally linked; for example, gig economy workers tend to have worse health outcomes than other workers because they attend fewer healthcare appointments due to the fact that they are not entitled to paid time off for healthcare appointments and cannot afford unpaid time off. Here, the only conventional kinds are ‘employed workers’ and ‘self-employed workers’, but these kinds scaffold social practices which give rise to a further explanatory kind, ‘gig economy workers’.

I then show how this practice-focused account can usefully be applied to the project of understanding and eradicating oppression, taking as my example of the oppression of trans women. Each of the following claims have been made by transfeminists and other proponents of trans liberation: trans women are oppressed as trans people, trans women are oppressed as women, and trans women are oppressed in virtue of not being counted/treated as women. The relationship between the claims has generated some confusion, and each claim has been disputed. I demonstrate that the practice-focused model of social ontology can precisely identify (1) the differences between these claims and (2) the information that is needed to support each claim (which turns out to be easily available in all three cases).

Finally, I also use this case to show how the practice-focused model sheds new light on an important and too-often overlooked insight: that emancipatory theorizing is itself a social practice, and as such is not separate from its own object of investigation (cf. Horkheimer [1937] 1972). In the terms of the practice-focused model, theorists investigating oppressive social ontologies can, through the activity of theorizing, affect the framing and scaffolding of those ontologies, for better or worse. Given the complexities of oppressive ontologies, this generates significant hazards to which emancipatory social theorists must be alert.

Culture Does Not Explain Behavior

Todd Jones—*Culture Does Not Explain Behavior*

University of Nevada, Las Vegas

One of the ways in which some social scientists and ordinary speakers seek to explain certain behaviors of group members is by describing how the behaviors were caused by the culture of the group. In this paper, I will argue that it is best not to think of culture as explanatory in this way.

At the heart of impulse to explain certain behaviors in terms of culture is a sound idea: When social scientists and others say that an act was done because of culture, they usually mean there is a disjunctive set of social factors present, any one of which would have caused that result in the right circumstances. Many of the features of the actual causal chain involved in producing this event were not strictly necessary, for if this set of circumstances had not produced this behavior (e.g. prayers before dinner), a related set would have. We talk about “culture” (and cognate terms like “norm” and “custom”) producing certain results, when we think that various different social and psychological processes all could have produced that result. Typically, we use “culture” to explain actions done because of various types of social pressure, or imitation. We consider culture to be the cause because none of the more particular causes is counterfactually necessary to produce actions of this type.

Despite this, cultural explanations are both overly broad and overly narrow. They are overly broad because more specific explanations of events in terms of particular environmental inputs and psychological states are always more precise and accurate than vague disjunctive cultural accounts. If what one is explaining, however, is a broader class of similar actions, a broader disjunctive account may well be appropriate. But in cases where this is so, there are usually many more factors besides social pressure or imitation that are likely to lead members of group X to do the behaviors in the Y set— so that a general disjunctive “cultural” explanation is actually not disjunctive enough. We should explain such events in terms of a wider set of disjuncts than the prototypical “cultural” forces.

Does the poverty of explanations in terms of culture mean that the entire discipline of Cultural Anthropology rests on a mistake? It does not. Although there are some exceptions, a careful look at anthropological explanations throughout the last century shows that what anthropologists of numerous different schools have usually sought to do is explain the items in a culture: the widespread presence of certain beliefs, practices, or artifact types in a particular group. Anthropologist rarely seek to explain things with culture. How we should best explain practices that are local to particular groups is an important and difficult question. But we shouldn't explain them with culture.

The Ontological Identity of the Family: Sociality, Purpose, and Care

Laura Kane—*The Ontological Identity of the Family: Sociality, Purpose, and Care*
University of Tampa

It is challenging to identify a conception of the family that is uncontroversial, uncontested, or inclusive of the many diverse groups who understand themselves as such. Consequently, there are many different interpretations of what the family should be – its desired member composition, its primary purpose, and its cultural significance – and many different examples of what families actually look like across the globe. In this paper, I propose that we understand the family as a unique social group with a particular primary purpose – to provide care in intimate settings for the mutual flourishing of all family members. The important point to note here is that this primary purpose of the family is what sets it apart from other social groups: the intimate caring relations found within the family are not found in other kinds of social groups.

Social groups are comprised of members who knowingly share a common feature with one another – a belief, a value, a practice, etc. – that differentiates one social group from another (Gilbert 1989). Like many social groups, a family is comprised of a certain number of group members who knowingly share some common feature with one another. This common feature I take to be the primary purpose that establishes the social group; I argue that we may determine what differentiates social groups only by identifying the primary purpose that guides each of them. Identifying the purpose of a social group brings the shared beliefs, values, practices, etc. to the forefront of group identification. I argue that a family qua social group should be recognized as having a particular primary purpose that originates from the family and for the sake of the family. By this I mean that the primary purpose of the family ought to be something that only pertains to the functioning of the family as a social group independent of all other purposes that the family might serve in different contexts.

I identify the primary purpose that the family serves for its members by identifying the normative criteria for members of a family to commit to certain kinds of goals and activities. I examine the most paradigmatic conceptions of the family that are based upon the supposed primary purpose that the family serves for its members and for the state, and ultimately conclude that the primary purposes that ground these accounts are not sufficient for distinguishing the family from other social groups; hence, they do not capture the unique primary purpose of the family.

I propose that the primary purpose of the family is to provide care in intimate settings for the mutual flourishing of all family members. This proposal calls for us to conceptualize the family as a social group committed to the goal of maintaining the well-being of all family members and the family unit as a whole through joint activity that works toward this shared goal. This joint activity is comprised of the practice of active care that family members are obligated to perform for one another by virtue of their joint commitment.

How to Make a Rule: A Defense of a Supposedly Refuted Theory

Jeffrey Kaplan—How to Make a Rule: A Defense of a Supposedly Refuted Theory
UC Berkeley

Human-created rules, often called "social rules," are plausibly at the heart of much of social reality. At minimum, they are at the heart of the existence of normative practices, such as fashion, etiquette, games, and law. What are the conditions for their existence? An old answer comes from H.L.A. Hart. His theory of law rests on a theory of rules which came to be called "The practice theory of rules". The practice theory offers two sufficient conditions for the existence of social rules.

The problem is that the practice theory is nearly universally rejected. It faces two main objections. The first objection dates back to G.J. Warnock (1971), and has been cited and relied on by dozens of philosophers of law including Scott Shapiro and Andrei Marmor. The objection comes in the form of a counterexample - a case in which the two conditions set out by the practice theory are met, but in which there intuitively is no social rule. As far as I am aware, every philosopher who discusses the counterexample accepts it as decisive.

The second objection is newer and perhaps gets at a deeper problem with the practice theory. The two conditions offered by the practice theory are descriptive. Rules, however, and the legal systems built out of them are thought to be normative. The descriptive explanatory resources of the practice theory are thought to be inadequate to explain a normative phenomenon.

In this paper, I offer solutions to both problems and argue that the practice theory has not been refuted. In response to the Warnock counterexample, I argue that there is a variety of rule that has been ignored. We see this by revisiting an old, but underappreciated distinction made by Rawls (1955). In response to the normativity problem, I argue that though mention of the normativity of law is now ubiquitous, it is widely misunderstood. Once we are clear on the sense in which law is normative, we can see that Hart's practice theory has no problem accounting for it.

The upshot of this is not only that an old theory of social entities is still plausible - though that would be interesting enough on its own. More broadly, understanding the nature of the kind of normativity exhibited by practices as diverse as fashion, etiquette, games, and law we see that it is, at least in principle, possible to explain these social practices by appeal exclusively to descriptive states of affairs.

Perduring Languages

Jessica Keiser—Perduring Languages

The Australian National University

There is a famous position about the metaphysics of ordinary objects—advocated by David Lewis and Ted Sider, among others—which takes them to be perduring, four-dimensional objects. Sider suggests that this position is motivated by its ability to diffuse classical puzzles about the identity of objects over time. I propose that an analogous view about languages—which takes them to be perduring object--can be motivated by its ability to solve a seemingly intractable puzzle about the nature of language; namely, how to reconcile our conception of language as an abstract semantic system with our conception of language as a changing social construct. (David Kaplan’s metaphysics of words has also been likened to four-dimensionalism about ordinary objects, though I use a completely different strategy which avoids the issues plaguing his account.) According to this view, languages can be modeled as series of functions from sets of possible utterances to functions from contexts to meanings. They are connected to linguistic populations (also modeled using a perdurantist, set-theoretic framework) by conventions-at-a-time. This proposal manages to do justice to the empirical work of linguists who model languages as static, abstract semantic systems, while at the same time providing the tools to accommodate real-world linguistic phenomena such as change, contingency, context-sensitivity, dialectical variation, gerrymandered acoustic and orthographic properties, and complex and fluid speech communities. It is further motivated by the fact that it provides a complete and reductive account of the metaphysics of expressions and their identity conditions, as well as the relation between word types and tokens--in contrast to previous theories which tended to provide partial or non-reductive accounts.

What Is the Relevance of Social Ontology for Social Scientific Research? A Comparison of Little's and Epstein's Conceptions of Social Ontology

Jo-Jo Koo—What Is the Relevance of Social Ontology for Social Scientific Research?
Grinnell College

Brian Epstein (esp. 2015) and Daniel Little (esp. 2016) have each recently put forward eye-opening accounts of what social ontology (broadly construed) has been and what it can and should be in the future. To begin with, Little argues extensively that social phenomena are fundamentally distinguished by their heterogeneity, ontological and causal complexity, plasticity, contingency, and path-dependency (Little 2016, esp. Ch. 1, 5-6). In particular, Little endorses the sort of ontological individualism that rejects reductive explanatory individualism and defends instead what he calls “methodological localism” (Little 2016: 75 and *passim*; cf. Little 2007). Roughly, this social ontology and its attendant philosophy of social explanation holds that while macro-level social events, facts, structures, and processes really exist and display relative explanatory autonomy in relation to the attitudes and activities of individual agents, the causal efficacy of macro-level social phenomena must always supervene microfoundationally on the activities of meso-level entities and micro-level individuals that are inescapably socially constituted and socially structured. Little thus endorses the supervenience of macro-level social causes on the lower-level activities of meso-level entities or micro-level individuals (Little 2016: Ch. 3 and 4).

Now, Epstein has persuasively challenged the standard (or consensus) view of ontological individualism as a nonreductive supervenience thesis in contemporary social ontology (Epstein 2015: 33-49 and much of Part Two). His objection turns on showing how there can be social facts that obtain (e.g., the sudden bankruptcy of Starbucks Corporation due to an accidental power spike overnight that fatally damages its material properties, equipment, etc., thereby causing it to become insolvent the day after), despite the fact that there were no changes in any attitudes or activities on the part of the individuals or meso-level groups that compose the higher-level entities of which they are constituents during the time in question. Epstein's objection challenges, therefore, the rightness of construing ontological individualism as a nonreductive supervenience thesis. More generally, Epstein argues that all hitherto social ontologies, even the most plausible and subtle ones currently on offer, including presumably Little's, are still overly anthropocentric by wrongly holding that social facts are exhaustively determined by facts about the attitudes and activities of human individuals (Epstein 2015: 7, much of Part Two, and 279). In fact, Little explicitly acknowledges the persuasiveness of Epstein's objection here, but nonetheless claims (rather dismissively, it must be said) that the goodness of this objection fails to undermine the plausibility of methodological localism (Little 2016: 129-32).

A number of thought-provoking observations arise in this comparison concerning the relevance of social ontology for actual social scientific research. (1) Suppose Epstein is right about the shortcomings of ontological individualism as a nonreductive supervenience thesis. As some reviewers of Epstein's book have noted, Epstein actually provides very little discussion about what implications his recommendations about how to properly do social ontology have for actual social

scientific research (Risjord 2016: 137; Di Iorio and Herfeld 2018: 121f.; Little 2016: 132). This situation provokes the question of what is exactly the relevance of social ontology, understood as a (purely?) philosophical enterprise that aims to clarify the metaphysics of the social world, for the understanding and practice of actual social science. If the answer to this question is 'none' or 'very little', this would seem to reflect badly on the significance of social ontology outside of philosophy. The fact that Epstein has yet to elaborate these implications does not necessarily mean that they do not exist. But if he is committed to showing how a more adequate social ontology leads to doing better social science, which is after all one of his official pronouncements in his book (Epstein 2015: 7 and 278), he is obligated to show what this would look like, at least in outline. (2) By contrast, despite Little's explicit concession that Epstein has made a persuasive objection to ontological individualism as a nonreductive supervenience thesis, Little's methodological localism, to its great credit, is packed with many detailed considerations of actual social scientific research from different social sciences. There seems, then, to be a lot going in this respect for ontological individualism as a nonreductive supervenience thesis about the social. This state of affairs naturally provokes the following question: Is it more fruitful and preferable to endorse Little's social ontology and philosophy of social explanation, which cannot eliminate the nontrivial set of counterexamples that Epstein has put forward, but nevertheless succeeds in engaging productively and persuasively with many actual social scientific research? Or is it rather more reasonable to prefer a social ontology, or more precisely, a set of constraints and conceptual tools that Epstein has offered us about how to do social ontology properly (cf. also Epstein 2016), that seems to be argumentatively better supported and yet has little bearing (at least thus far) for actual social scientific research? (3) More generally in light of the above, what is really the relevance of social ontology for actual social scientific research? This talk will try to make some headway regarding the answers to these questions.

The Metaphysics of Intersectionality

Odin Kroeger—The Metaphysics of Intersectionality

University of Vienna

Dotson (2016) identifies multiple assumptions that, if true, would make it the case that inquiries into ‘race’ are more fundamental than those into ‘race’-and-gender, so that we could research racism simpliciter. Of course, this is unlikely. However, whether inquiries into ‘race’ are more fundamental than those into ‘race’-and-gender, needless to say, depends on the nature of those social kinds. So our metaphysics of social kinds should, at least, allow for one (or more) of those assumptions to be false—or, better yet, help us explain what makes it (them) false.

Yet, following Boyd (1991), social kinds are, typically, taken to be homeostatic property clusters. For example, Haslanger (2012, 2014a, 2014b, 2016) argues, simply put, that a kind, *K*, is social if and only if *K*-typical properties have a certain meaning in how we behave, to the effect that we ‘cluster’ those properties on the one hand and their meaning on the other into a category “*K*” and behave so that we make it the case that things that have those properties exemplify *K*. At first sight, this appears to strengthen the case that some inquiries are more fundamental than others. After all, if something exemplifies *K* and another kind *J*, then, it appears, I can inquire into *K* and *J* separately, and inquiries into *K* would seem to be more fundamental than those into *K*-and-*J**. As a consequence, there appears to be some tension between the insight that oppression is intersectional and the way we conceptualise the nature of social kinds.

However, drawing on Bright et al. (2015), who show how intersecting kinds can be modelled causally, I argue that social kinds are unified by the mechanisms that cause them to be exemplified, to the effect that even if two social kinds, for example, Black Woman and Black Man, share the same genus, their social meaning, their role-typical properties, that is, may still be quite different. This enables me to put forward a model that can explain why, for some kinds, the assumptions that Dotson identifies are false.

Discoverability and Contestability in Social Construction (Of Gender)

Nick Kröll and Bennett Helm—Discoverability and Contestability in Social Construction (Of Gender)

Franklin & Marshall College

According to social constructionist accounts of gender, gender is the product of some complex social fact that determines social roles, such that someone's gender is a matter of their occupying a particular gendered social role or, perhaps, identifying with such a role. This makes it difficult to see how the society itself could get those gender roles wrong by being mistaken about what existing genders actually are or by failing to construct alternative gender roles (as opposed to being wrong about matters of justice (e.g., Haslanger, 2012)). We propose and defend a framework that can help make sense of this possibility of the community getting gender wrong and hence of socially constructed facts as having greater objectivity than is commonly recognized.

Objectivity presupposes both **discoverability**---where objective phenomena are criterial for the correctness of one's experience---and **contestability**---whereby subjects can engage in substantive disagreements about those phenomena, partly on the basis of that experience. In science, particular experiences can ground the contestation of concepts only insofar as those concepts are embedded within and hence defined by a broader explanatory theory. For it is the background theory that rules out certain phenomena as impossible, so that experiences of such impossible phenomena provide reason to revise or reject that theory and its theoretical concepts.

We propose to understand the objectivity of socially constructed phenomena like gender in an analogous way. Gender concepts are embedded within and gain their content by a largely implicit partial "folk theory" (or ideology or ...) of how it is best for us in this society to live, a theory that prescribes and proscribes certain patterns of attention, responses, actions, attitudes, etc. Such a theory will thus rule out certain kinds of lives as impermissible, as not conducive to flourishing, such that experiences of such a life as actually or potentially flourishing provide reason to revise or reject that theory and its concepts. In particular, a folk theory of gender that ties gender closely to binary sex categories and so rules out transgender identities as incompatible with flourishing can be contested in light of contrary experiences, which provide reason to revise our existing concepts of **man** and **woman** and hence the theory itself. Likewise, a folk theory of gender that understands it in binary terms and thus rules out non-binary gender identities can be similarly contested, leading to the introduction of new gender concepts.

In the face of such experiential evidence and the reasons it provides, the resulting revisions to the society's gender concepts are intelligible as **improvements** in understanding gender, and allow members to discover new facts about themselves to which they were previously blind.

Collective Agents and Corporate Responsibility: Bringing Together Legal and Ontological Perspectives

Visa Kurki, Säde Hormio, and Pekka Mäkelä—Collective Agents and Corporate Responsibility University of Tampere, University of Helsinki, and University of Helsinki

What, exactly, are corporations and for what can and should they be held responsible? Corporations play a key role not only in modern societies, but also in various social sciences, including economics, political science, and sociology. In social science literature, including business ethics, there are various attempts to conceptualise the corporate world in terms of such notions as corporate citizenship or corporate social responsibility. Despite the multidisciplinary interest in corporations and their key role in shaping our realities, scientific discussion on corporate responsibility is lacking a unifying theoretical framework.

The literature on corporate responsibility suffers from its not being based on a proper, rigorous understanding of the nature of corporations, which combines the social and legal aspects of corporations. Philosophical accounts of corporate responsibility typically provide very sketchy accounts of the legal arrangements pertaining to corporations. Legally informed accounts, on the other hand, often lack philosophical sophistication. Our paper will attempt to cover both bases by offering a stringent analysis of the philosophical as well as the legal issues underlying corporate responsibility.

The paper will start by disambiguating a number of ways in which the term ‘corporation’ is used. Most importantly, the term may be used to refer to an organised and incorporated collective agent (a collective agent with a specific status) as well as to a kind of legal arrangement (the status itself). If one announces the plan to found a one-person corporation, one is referring to a legal arrangement; if one talks of the great company culture of a certain corporation, the focus is on the organisation. Many contemporary authors have not paid sufficient attention to this distinction. The paper will focus on corporations in the former sense. Such corporations are simultaneously socially and legally constituted. They are socially constituted in that they consist of people acting together toward a common goal. However, corporations are also legally constituted in that they can only exist within the context of a legal system.

Our paper attempts to provide a stringent analysis of what exactly distinguishes incorporated collective agents from nonincorporated collective agents. Many of the traditional accounts prove simplistic: for instance, the popular claim that corporations are group agents with an institutional status, though not incorrect, fails to distinguish corporations from unincorporated business entities such as partnerships. The paper will seek to provide a more nuanced analysis of corporations as well as explore certain implications of this analysis.

Is Face an Institution?

René Lacroix—Is Face an Institution?

University Paris 3 & the French National Center for Scientific Research

Face is one of the central themes in Erving Goffman's writings (see, among others, 1967b). It can be described as the positive or negative image that others form of us in a given situation on the basis of our behavior. For instance, in certain contexts, stumbling, having a stain on one's shirt, misspelling a word or not understanding a joke trigger a negative evaluation of us by those who witness our behavior.

In this talk, I will use Searle's (1995, 2005, 2010) theory of institutional facts and Epstein's (2015) criticism of it to analyze the phenomenon of face. I will argue that face, which is a social phenomenon, is more specifically an institution in Searle's sense of the term, though of a particular type. As institutional facts, face is constituted by a rule of the kind X counts as Y in C: in context C, stumbling (X) counts as losing face (Y). It is a constitutive, not a regulative, rule: face is not a preexisting phenomenon that is regulated by the rule; rather, it comes into being by the very existence of the rule. Thus, it is observer-relative.

A trickier question is whether face implies deontic powers. Losing face does not give any official rights or obligations in the same sense as being an attorney or owning a property do. However, it does impose an obligation on the participants, including the one who has lost face: they now have to restore the "ceremonial order" (Goffman 1967c) in the ongoing interaction by means of a corrective process (Goffman 1967b: 19-23). As in the case of honorific statuses, the deontic powers involved here are only "weak" or "limited" ones (Searle 2010: 24). In this sense, face is not a prototypical institutional fact.

According to Searle (1995: 97), the Y status in his formula can be imposed on three ontological categories of phenomena: people, objects and events. Face-losing clearly pertains to the first category. A more complex question is that of the facts that ground face (Epstein 2015, 2016). These facts may be located in/on someone's body (movements, clothes, odors, speech, etc.), or they may be an accompanying object or person (e.g. my child, sitting next to me in the bus). Other possible ways to classify the facts that ground face will be suggested.

As we will see, face may be involved in iterated structures, as do other institutional facts (Searle 1995: 80). This happens when someone loses face in virtue of being in a certain situation (a job interview, for instance), the situation itself having come into existence as a result of a constitutive rule.

Climate Change and the Collective Responsibility for Infrastructure

Lauri Lahikainen—*Climate Change and the Collective Responsibility for Climate Change*
University of Tampere

Many climate ethicists have argued that the responsibility for the mitigation of and adaptation for climate change is collective in nature. There are many ways to conceptualize and account for this collective responsibility. We can be responsible for climate change as members of culpable collectivities or quasi-collectivities such as the global elite, the North, the polluters, and so on. On the other hand, collective responsibility may stem from ability of certain collectives to act. Thirdly, it may be that we have a duty to form collectives that have the power to address climate change if no such collectives exist yet. Following Epting's (2016) suggestion that there are moral dimensions to infrastructure, I argue that we will understand these forms of collective responsibility better if we focus on how the relevant collectives are related to both fossil fuel powered and greenhouse gas producing infrastructures. The culpable or prospectively responsible collectives are then not defined by their emissions per capita or some other metric focusing on emissions, but rather by how much these collectives have historically had or now have power over what kinds of infrastructures are built and maintained. As users and beneficiaries of shared but unequally accessible infrastructures, we may also have a degree of responsibility for climate change. The responsibility to form capable collectives becomes relevant if existing institutions are unable or unwilling to transform current systems of infrastructures. Besides responsibility for climate change, I suggest that we may have a more general collective responsibility for the infrastructures that we build and use, but this responsibility may be difficult to see or realize due to current social arrangements and the nature of infrastructure itself as partly phenomenologically non-transparent. Finally, infrastructures are an integral part of our social arrangements, and thus new kinds of infrastructures may make new kinds of collectives possible.

Social Wrongs

Arto Laitinen—Social Wrongs

University of Tampere

In this paper we elucidate the notion of “social wrongs”. We try to show how that idea differs from moral wrongness on the one hand, and how the social wrongs that are social pathologies differ from the wrongs discussed in mainstream political philosophy (such as illegitimacy) on the other hand. Both social and political wrongs share a feature with natural badness or wrongness (illnesses of organisms) as well as malfunctioning artefacts or dysfunctional organizations: they violate so called ought-to-be norms; they are not as they ought to be. In contrast, moral wrongness violates ought-to-do –norms.

In Section One we go through various varieties of wrongness: monadic (1.1) and dyadic wronging (1.2), and being wronged by practices, institutions or structures (1.3) In Section Two the natural and artificial cases of there being something wrong with an organism or a system (2.1), the narrowly political wrongs of systems of governance (2.2), and the broader category of social wrongs are discussed. (2.3).

Section Three will address the question whether the applicability of the category of social wrongs presupposes a social ontology in which social reality is normatively constituted (e.g. as practices presupposing norms and participatory understandings; or as social arrangements having a function which can then malfunction).

Collective Responsibility to Prevent AI Systems from Becoming Responsible

Arto Laitinen and Joanna Bryson—Collective Responsibility to Prevent AI Systems from Becoming Responsible

University of Bath and the University of Tampere

This paper starts by noting that two kinds of arguments exist against AI systems being responsible agents, or patients towards which moral agents have responsibilities. One argument claims that AI systems cannot have the relevant features – we simply cannot make AI systems that would be moral agents or patients. Another argument claims that we can build such systems but there are conclusively strong reasons against doing so (Bryson 2018).

We reject the first argument and accept that once AI systems meet the relevant conditions for moral agency, they are moral agents, and same for conditions for moral patiency. But we develop the second argument that we ought not make AI systems that meet those conditions.

What could be wrong with AI systems being responsible? Isn't responsibility a good thing for a system? We distinguish the categorical sense of fitness to be held responsible and the appraisal sense of something acting or functioning in a responsible (i.e. praiseworthy) manner as opposed to blameworthy manner.

We do of course accept that once an AI system meets the conditions for moral agency, it ought to act in a praiseworthy rather than blameworthy manner. The argument is against making AI systems that are fit to be held responsible, blamed or praised; or that are moral patients to whom others owe responsible treatment. We go through several reasons presented in the literature.

What this paper adds to that literature is first of all the question of distribution of responsibility: if there are things that ought not take place, who should see to it that they do not take place? All agents of course concerning their own actions, but we also argue for a collective responsibility: ultimately it is a matter of responsibility for the shape of societal ethics, which is everyone's shared responsibility.

Secondly, we address the worry of responsibility vacuums (or "gaps") created by "autonomous" AI systems. Would there not be an argument from vacuums to the need to make these AI systems responsible? And if the systems themselves are not responsible, who are? We argue that as there are plenty of responsibility-distributing techniques available (including strict liability), such worrisome responsibility vacuums are not necessary; social practices can be devised so that there is always some humans responsible in the forward-looking sense (even when no-one is to blame, and no backward-looking responsibility is available).

Two Paths to We (As I)

Dimitri Landa—Two Paths to We (As I)

NYU

Tuomela, Gold and Sugden, and others advance an account of joint intention defined by a "We-mode of reasoning. In this mode, there is an "instrumental [non-normative] group sense of collective commitment" whereby when agents jointly intend to x as members of the we-group, they think of themselves not as distinct individuals, but as the members of an x -intending we-group and implicitly bind themselves to doing what x requires.

If reasoning as We is to play a definitive role in a distinctly joint intention, it should sometimes produce a different effect than reasoning as I. Indeed, Tuomela (2006) and Gold and Sugden (2007) argue that, in the "We-mode," cooperation may be sustainable in the one-shot Prisoners' Dilemma (PD). The argument is as follows: if we would like (jointly intend) to accomplish a high-payoff outcome, and accomplishing that outcome requires that we play {cooperate, cooperate}, then each of us individually does our part(s) to bring that about -- that is, chooses to cooperate. In effect, it means that all further appeals to individual rationality are superseded by appeals to instrumental action on behalf of We as a single unitary agent.

Despite the enthusiasm about the "We-mode," however, there has been little in the way of developing an account for how individuals reasoning as I transition to reasoning as We, and there are good reasons to view such a transition as an important puzzle for a theory of social ontology that puts a premium on respecting basic commitments to actor rationality. I develop an account of what such a transition may entail. At the core of the account is the idea that individuals may experience empathy-inducing events that affect their perception of the nature of the situation they encounter and of other individuals in it. I argue that such empathy-inducing events can lead to two distinct paths to reasoning as We. The first path is through an increase in sympathy with others. As a consequence of such an increase, their welfare may matter to me to the point that my objective function becomes effectively defined on behalf of Us, rather than of me alone. The second path is through "institution-utilitarian" (Parfit 1984; Hardin 1988) reasoning that runs as follows: "we (I and others) would like to accomplish x , therefore we must seek what will make x an equilibrium implementable choice."

Critically, while both of these paths lead to reasoning on behalf of We, it would be a mistake to see either of them as leading to a preference to cooperate in a one-shot PD. The effect of the first path is to change the underlying game away from PD to a different game in which players' payoffs are increasing in each other's. In that game, there may be a "cooperate" action that has the same substantive interpretation as "cooperate" in the original PD, and so a project important to Us may now be realized through cooperation when it was not before. But this should be properly seen as the effect of a change in the game, not of how the original game is played.

The effect of the second path is to induce a search for how to transform the original game into one in which a mutually important project x could be realized through individually rational play. One common way of doing this is to introduce a schedule of sanctions that alters the payoffs in the underlying game, either to eliminate undesirable equilibria or to create an attractive equilibrium *de novo* (Almendares and Landa 2007) that enables the parties to sustain a "cooperative" play that they could not in a one-shot PD. Ironically, though, the joint resolve to transform the original game

in this fashion -- which corresponds to individually rational commitments -- does not make sense if we are to follow through on the "We-mode" interpretation of joint intention in the context of the one-shot Prisoners' Dilemma. Both Gold and Sugden (2007) and Tuomela (2006) appear to endorse the view that the individual commitment to "We" means setting aside any further appeals to the individual rationality of action. Gold and Sugden argue, in particular, "If I am to reason instrumentally, I cannot simultaneously think of myself both as a unit of agency in my own right and as part of a unit of agency which includes you" (p. 14). The institution-utilitarian approach puts this in doubt. Actions taken to bring about an equilibrium-implementable Pareto-superior outcome appear simply to fall outside the purview of the "We mode" view.

On the account I develop, then, (1) reasoning as We can be explained as an important social phenomenon without committing oneself to metaphysical shifts in actor identity; and (2) reasoning as We does, indeed, sometimes produces a different effect than reasoning as I, but the ways in which this comes about do not require surrender of commitments to individual rationality.

Scientific Consensus and the Discursive Dilemma

Helen Lauer—Scientific Consensus and the Discursive Dilemma
University of Dar es Salaam

Nowadays, scientists successfully committed to public service work within vast income generating, institution-dependent, tightly orchestrated, highly controlled collaborating networks designed to satisfy a range of expectations that are often vehemently opposed. For example the obligation of epidemiologists in the global health arena to pursue standards of maximal accuracy and optimal precision on the one hand, pulls against their simultaneous obligations to err on the side of caution, to provide decisiveness of judgments and advice under conditions of extreme uncertainty, in order to preserve public confidence and safety, to satisfy industrial investors and further government agendas (Lauer & Shenton 2017).

An authorized discourse represents the large professional research webs managing such potential conflicts of epistemic interest, dedicated to producing the global impression of a stable scientific consensus. Oversimplified claims are promulgated on behalf of these large conglomerates that starkly deviate from the judgments and viewpoints of the individual scientists and the research sub-groups who constitute their membership. Thus an apparent paradox emerges in (i) the exhibition of doctrinal rule-compliance versus practical norm-following of group behaviour defining what it means for collective behaviour to be scientific (Ludwig 2017b) and also in (ii) the attribution of judgments aggregated as collective scientific consensus versus attribution of judgments as distributed among the individual members contributing to that consensus.

To account for both of these divergences (i) and (ii) above, I follow a suggestion of Ludwig (2017a) in applying a startling theoretical result about judgment aggregation labelled the discursive dilemma, authored by List (2012) and corroborated by other social choice theorists. For example, when research collectives compromise maximal diversity, inclusiveness, impartiality, and rigor in the management of data and explanatory conjecture, their concession to non-epistemic utilities reflect the 'relaxation' of more general conditions identifying ideally rational judgment aggregation procedures, i.e. universality, anonymity, and systematicity of procedure (List 2012, Pettit & List 2002). The discursive dilemma result shows that some sort of concession is necessary to avoid inconsistencies between research hives' judgments and the official consensus standing for collective epistemic choices of epidemiologists as contributors to a current scientific genre.

The apparatus provided by List and by Pettit might efficiently supplant two alternative accounts of the inconsistency between distributed and collective intentional states: one model posits non-epistemic priorities undermining scientific excellence (e.g. Kitcher 1993). A second model imbues institutionalised science with epistemic agency irreducible to the rational processes accorded to its individual members (Pettit 2017). These popular approaches are accused of being ideologically and ontologically effusive, respectively.

Additionally, the discursive dilemma result may help determine whether there are calculable limits to the disparate aims that large scale scientific collaborations can rationally serve.

Collective Responsibility Revisited

Holly Lawford-Smith—Collective Responsibility Revisited
University of Melbourne

It has been standard to assume that if an agent causes harm by way of an action for which she is not excused, then she is responsible for that harm. In this paper I will argue that this is not true: the exercise of collective agency (including collective moral agency) can produce harms for which the collective agent is not excused, and it can still be the case that the collective agent is not responsible for the harms. The argument proceeds in two parts. In the first part, I argue that individual agency (including moral agency) is not enough to justify individual responsibility in specific class of cases. There are cases involving what I have elsewhere (with co-authors) called the 'Many Times Problem', in which individual acts performed by the same agent at different times cannot be unified in a way that justifies responsibility for the outcomes of the acts taken cumulatively. In the second part, I argue that this shows that the idea of collective responsibility needs to be revisited, at least for this class of cases. It's not enough that collectives are agents (including moral agents). Instead, the actions they perform must be unified in a particular way across time. I argue that we can hack John Broome's discussion of discounting the future, and Derek Parfit's discussion of psychological continuity, to get a plausible account of how collective agents' actions must be unified in order for there to be responsibility for their outcomes across time.

Collective Moral Luck

*Edward A. Lenzo—Collective Moral Luck
University of Memphis*

According to what is known as the “Control Principle” (CP): We are morally assessable only to the extent that what we are assessed for depends on factors under our control. However, Thomas Nagel asserts, factors outside of our control are in fact morally significant, and make a difference to our ordinary moral judgements. Nagel formulates the notion of moral luck accordingly: “where a significant aspect of what someone does depends on factors beyond his control, yet we continue to treat him in that respect as an object of moral judgement, it can be called moral luck.”

There is a tension between CP and moral luck: Nagel rightly notes that the consequences of our actions, the circumstances in which we find ourselves, and even who we are as people, always (or almost always) depend on factors outside of our control. Moral luck becomes a problem: if factors outside of our control are relevant to our ordinary moral judgements, and if our actions are nearly always subject to such factors, then, if we maintain CP, almost none of our ordinary moral judgements are left intact. Nagel also believes that we cannot simply give up CP, as its application is legitimately persuasive of the importance of the absence of control for moral assessment.

Nagel and Bernard Williams – often cited as coining and developing the concept of moral luck – both agree that our constitution (i.e., who we are as people, our dispositions, “inclinations, capacities, and temperament”) are “as conditioned as anything else.” In other words, even who we are is subject to moral luck, and thus there is a (currently unsolved) puzzle as to whether or not we can be held morally responsible even for our own will.

The problem of moral luck is almost entirely couched in individualistic terms: it is the individual human agent who is considered to be the subject of moral luck. Assuming certain types of collectives are moral agents (ala French – and List and Pettit), I ask: are these collectives subject to the problem of moral luck? I here focus on the constitutional notion of moral luck, its application to collective cases, and implications for collective responsibility.

In section I, I introduce and motivate the general problem of moral luck. In section II, I explicate the specifically constitutional notion of moral luck as regards individual moral agents, then apply this notion to the collective case, arguing that collectives are far less susceptible to constitutional luck than are individuals (even those of which the collective consists), since they have a high degree of control over their own constitutions; they are intentionally self-constituting. In section III, I conclude that collectives can therefore be held responsible for their own constitutions and actions in a wide number of cases the individualistic analogues of which would cause us to excuse or puzzle over the unlucky agent.

Mutual Responsiveness in Action and Shared Action Spaces

Felipe León—Mutual Responsiveness in Action and Shared Action Spaces
University of Copenhagen

According to a leading philosophical account of small-scale cooperative joint action (Bratman 1999, 2014), central instances of it require from participants to be mutually responsive in action to one another, i.e. they demand the flexible tracking and adjustment of an agent's actions to the actions of co-agents during joint action execution. This leading account has been criticized for not saying much about the 'body glue' that mutual responsiveness in action apparently requires, at least in cases that involve the spatio-temporal proximity of the involved agents (Pacherie 2015, p. 24). On the face of it, there seems to be a certain puzzle about the relationship between embodiment and mutual responsiveness in action. In normal circumstances, an individual agent receives an on-going flow of information about his or her own body, in particular (although not only) during action execution. The puzzle arises because certain activities involving mutual responsiveness in action, such as lifting and carrying a heavy object together with someone, seem to demand that each agent receives and integrates information about a co-agent's embodied perspective in a fluid and flexible way. The question of how this might be possible hasn't gone unnoticed in the literature, particularly in psychological research on joint action. A number of cognitive mechanisms that plausibly enable joint action execution have been investigated, including co-representation of tasks, co-representation of perceptions, joint attention, and alignment mechanisms (Pacherie 2015, Knoblich et al. 2011, Gallotti et al. 2017).

The aim of my presentation is to investigate the proposal that joint action execution is supported by "shared action spaces" (Pezzulo et al. 2013; Pacherie 2015). I distinguish two ways of interpreting this notion, and argue in favour of one of them. According to the first interpretation, a shared action space results from each agent supplementing his or her individual action space with information about another agent, via for example perspective-taking and simulation. These and other cognitive mechanisms would provide each agent with the necessary information about another agent's individual action space in order for joint action execution to get off the ground. According to the second interpretation, the availability of a shared action space does not depend on a supplementation of each individual's action space, but rather on the social re-framing or recalibration of his or her action space as a second-personal or I-you space, given the suitable presence and readiness to engage of other agents. This second interpretation of a shared action space provides for an attractive alternative to the cognitive demandingness that the first interpretation seems to involve. It can accommodate the relevance of perspective-taking and simulation, by pointing out that they may take place on the basis of a shared action space (as an I-you space). And it accommodates the idea that a shared action space may include distinctive "social affordances", i.e. joint action possibilities that are not available to one individual in isolation (Pezzulo et al. 2013, p. 2).

What Should We Do? The Bloomington School and the Citizen's Question

*Peter Levine—What Should We Do? The Bloomington School and the Citizen's Question
Tufts University*

Normative political theory typically asks: "How should things be?" or "What should be done?" For instance, to combat global warming, carbon should be taxed in every nation. However, for an active citizen who has limited influence and information, the question is much more difficult. Citizens must ask "What should WE do?" where "we" refers to a concrete group of actors large enough to make a difference but small enough for the individual agent to have influence.

"What should we do?" raises generic theoretical issues: how to overcome collective-action problems, how to deliberate about values when people disagree, how to confront opponents outside the group, how to define complicity and loyalty, and how to relate small-scale activity to large-scale problems.

I argue that Elinor Ostrom and her colleagues have made essential contributions to addressing the citizen's question. Their research reflects a Copernican turn for political theory: from society to citizens. However, Ostrom's Bloomington School focuses on cases in which ends are relatively clear and the problems involve coordination. Their work is therefore less useful for situations in which people disagree about ends, or in which some people's ends are unjust or unwise. It is also less useful for situations in which an intransigent power stands in the way. For these situations, I argue, the deliberative democratic tradition exemplified by Jürgen Habermas and the nonviolent social movement tradition exemplified by Gandhi and King offer complementary insights. Combining these traditions is the way forward.

Belief and Linguistic Understanding

Ronald Loeffler—*Belief and Linguistic Understanding*
Grand Valley State University

According to standard approaches to belief, believing that *p* – where *p* is a non-social first-order content, such as that it is raining – is intrinsically private, in the sense that forming the belief *per se* involves no sense on the subject's part that there are, or may be, other epistemic subjects who might share or critique the belief. Of course, believers usually do have such a sense of the actual or potential presence of other believers. However, on the standard view, having this sense is extrinsic to believing that *p* itself, in that it requires something extra – such as higher-order thoughts about other believers or off-line mental simulation of their doxastic perspective.

Drawing on mode accounts of collective intentionality, as offered by Raimo Tuomela or Hans Bernhard Schmid, which treat phenomena such as joined intentions to act as members of a group as rooted in cognitive attitudes (as opposed to cognitive content) I propose to treat belief similarly, as a cognitive state that qua attitude (or mode) is intrinsically intersubjective. Taking seriously that believing that *p* is taking a propositional attitude of holding true towards a proposition or sentence that *p*, and regarding truth as objective in the Kantian sense of true for every rational being as such, I propose that, just by taking the belief attitude towards the content or sentence that *p*, the subject *ipso facto* treats *p* as to be endorsed (held true) by everybody. Believing that *p* is thus intrinsically intersubjective in that it involves an implicit normative expectation by the subject that everyone should share the belief, rooted in the belief attitude taken towards *p* itself. This normative expectation is implicit in that, because it is a matter of the belief qua attitude, not of content, it is not a matter of applying any concepts – such as the concepts of belief or truth, or related normative concepts – and a fortiori not a matter of higher-order thinking, nor a matter mental simulation.

I then relate this proposal to contemporary neo-pragmatist efforts, such as Robert Brandom's, to explain linguistic understanding in terms of implicit mutual normative expectations between speaker and interlocutor in linguistic communication. I argue that if we, first, treat beliefs as attitudes towards sentences (rather than propositions); second, treat linguistic understanding as a matter of normgoverned inferential role (the role of the sentences in good reasoning); third, treat the attitude of taking an inference as good analogously to the belief attitude (as an attitude taken towards pairs, triples, etc. of sentences – rather than individual sentences, as in the case of belief – that intrinsically involves an implicit normative expectation on everybody to treat the inference involving these pairs, triples, etc. as good) – a picture emerges of linguistic communication as shot through with implicit, mutual normative expectations, concerning what we should believe and how we should reason, and of linguistic understanding as an aspect of such implicit mutual social recognition.

Conventions and Status Functions

*Kirk Ludwig and Marija Jankovic—Conventions and Status Functions
Indiana University and Bloomington & Davidson College*

Conventions and status functions are central features of social life. How are they related? In this paper, we argue that

- a) there is a variety of convention that has not been adequately identified in the literature,
- b) status functions constitutively involve this variety of convention, and
- c) what is special about it explains the central feature of status functions, namely, that objects with status functions can perform their functions only insofar as they have been collectively accepted as having them.

We develop this project as follows. First, we describe a social kind we call an effective agreement by appeal to cases in which people settle into a solution to a coordination problem, e.g., about when to arrive at a café for a book club meeting, without an explicit agreement. We argue that this involves members of the group acquiring interlocking conditional we-intentions with respect to a collective action plan that solves the problem. Specifically:

Members of G are in an effective agreement to act in accordance with a collective action plan P , in circumstances C involving some members of the group G , if and only if:

- (i) the members of G collectively intend that whenever any of them are in a circumstance of type C , they are to achieve end E ;
- (ii) each member of G we-intends to act in accordance with P in order to achieve end E , whenever she is in a circumstance of type C ;
- (iii) there is an alternative, P' , to P in C by which the members of G could collectively achieve E in any instance of C by acting together intentionally in accordance with P' .

Second, we argue that this meets the pre-theoretic desiderata for being a kind of convention: it is arbitrary, social, stable, and reciprocal.

Next, we argue that it is an important but overlooked social kind by showing its central feature: the imposition of status functions. A key feature of status functions (e.g. \$20 bill or the queen in chess) is that objects with status functions can perform their functions only if they have been collectively accepted as having them. We explain this by identifying the functions involved as defined by constitutive rules for essentially intentional collective action types that leave open what things are to play certain roles in the activity. This presents a coordination problem for agents who intend to instantiate the action type. It requires them to coordinate on the same objects in the relevant roles. The underlying attitude structure is exactly that exhibited in effective agreement. This explains what collective acceptance comes to (in a new way) and why it is required for something to have a status function. Finally, we contrast this concept of a convention with Lewis conventions, which require both more (e.g., common knowledge), and less (e.g., no shared intentions) than the kind we identify. We do not reject Lewis conventions, but they are not fitted to explain what is special about status functions, and so, we argue, an important concept of convention has been overlooked.

Shared Agency and Mutual Accountability

*Taylor Madigan—Shared Agency and Mutual Accountability
Stanford University*

'How must individuals be related to each other in order for those individuals to count as sharing an intention?' The reductionist answer maintains that a shared intention can be constructed using the building blocks of the attitudes of the group members (e.g., Bratman 1993, Tuomela and Miller 1988). According to this view, each participant in the group activity must have, in virtue of being a participant, some set of attitudes that are appropriately related and responsive to the attitudes of the other participants, and together these attitudes constitute the shared intention. This offers the prospect of a theoretically parsimonious account of shared agency, in the sense that it employs the conceptual, metaphysical, and normative resources that have been extensively studied in the philosophy of action.

However, the reductive approach faces serious obstacles. The challenge I address is that participants in shared agency seem to be bound together by directed obligations and corresponding entitlements to hold each other accountable for doing their respective parts in the group activity. For example, suppose you and I go on a walk together, and part way through the walk I abruptly stop and leave without explanation. It would seem that you are entitled to rebuke me for my behavior, which suggests that I have an obligation to do my part, and obligation that I am not living up to (Gilbert 2009). Theories that reduce shared intentions to interpersonal structures of personal intentions seem unable, in principle, to account for the mutual accountability distinctive of shared agency, because personal intentions are not sufficient to underwrite directed obligations and entitlements.

I present a reductive proposal that vindicates Gilbert's insight about the importance of mutual accountability without relinquishing the methodological virtues of reductive accounts. The key move is to adapt the Searle-ian notion of a status, which I argue can explain mutual accountability and be reduced to interpersonal structures of personal intentions. For example, if you and I are on a date, we each intend that each has the status date, which specifies certain ways in which we take each other to be accountable to one another. In turn, for an agent to have a certain status is explained in terms of a relevant set of participants intending that she have the status, which is to say that they intend that she be treated in certain characteristic ways. For example, the members of a club can bestow the status of club president to a certain subject, which she has simply in virtue of the club's members being disposed to intentionally react to her in certain ways in certain circumstances.

Hence, I argue that statuses are reducible to interpersonal structures of personal intentions and thus are a flexible yet powerful tool at the reductionist's disposal, which she may employ to explain mutual accountability without introducing new and irreducible normative resources such as "joint commitment."

Hybrid Agents and Responsibility Sharing

Pekka Mäkelä and Raul Hakli—Hybrid Agents and Responsibility Sharing

University of Helsinki

We study whether humans and technological artefacts, such as robots, can form hybrid agents that would be fit to be held morally responsible for their actions. Several arguments for this possibility have been presented but we argue that they have not been successful. We identify three argument forms that have been employed to argue for shared responsibility between humans and machines: (i) Argument from gradual properties, (ii) Argument from responsibility gaps, and (iii) Argument from extended agency. We analyse these arguments and aim to show that they are invalid.

Dangerous Speech and Collective Action

Teresa Marques—Dangerous Speech and Collective Action
University of Barcelona

Recent years saw an upsurge of populist discourse in Europe and in the US. This is manifest in the reappearance of xenophobic, racist, or nationalistic speech. This upsurge is correlated with violence. There are 900 hate groups active in the US, and there have been at least 100 Alt-Right related killings (SPLC data). Understanding the mechanisms that correlate speech with discrimination and violence is a necessary step to explore the most effective forms of counterspeech. I will use populist and dangerous speech as paradigms of evaluative motivational discourse. A theory of evaluative motivational speech should explain (a) how it expresses affective and action-guiding attitudes, (b) how it builds a connection among the people who accept it, and (c) how accepting such speech commits the audience to do their part in future actions. These desiderata become perspicuously clear in extreme cases of bigotry formation (Jeshion (2016)), of expressions of exclusionary pride (Snyder (2017)), or of dangerous speech (Maynard and Benesch (2016)): speech acts whose force is capable of encouraging approval of violence by the audience. How do words lead to action? In this paper, I will argue that to answer this question we must rethink the role of expressive presuppositions, and the notions of common ground and conversational score. After Bernard Williams (1980), I propose that a conversation's score (after Lewis 1979) and a conversation's common ground (after Stalnaker 1973, 1974) must include more than propositions that are accepted as true; they must include a motivational set that includes shared plans, evaluative dispositions, and norms. This offers an explanation of how the acceptance of evaluative speech by an audience can lead to attitude and behaviour modification. My proposed account can explain how acceptance of certain speech-acts compromises the audience to future courses of action. The account makes a distinction between the 'common motivational set' (which is part of common ground), i.e., those motivational attitudes that are actually common, and a 'motivational score', i.e. those attitudes that are part of the conversational score. I suggest that the notion of conversational motivation score can be thought after Bicchieri's 2016 work on social norms. I argue that it is useful to distinguish the conative attitudes and plans that are common, and those that people follow even though individuals may be indifferent or disapprove of aims, means, plans, or norms. Finally, and after Kutz 2000's minimalist account of collective action, I claim that the acceptance of an evaluative statement is an update to the motivational common ground, and an acceptance to do one's part towards common goals. Accepting problematic expressive presuppositions compromises an audience to participate in problematic actions. Finally, I suggest that Ayala and Vasylieva 2016's suggestion that audiences share responsibility in problematic contexts fits within the present account, and discuss how it applies to concrete cases of dangerous speech.

Biased Minds as Internalized Structural Phenomena: A Dual Account of Social Structure

Annette Martin—*Biased Minds as Internalized Structural Phenomena: A Dual Account of Social Structure*
New York University

There has been a great deal of debate about the relative importance of individualistic and structural remedies for social inequality. Both inside and outside of the academy, individualistic solutions have tended to focus on undoing implicit and explicit biases. In response, “structural prioritizers” like Haslanger, Anderson, and Ayala argue that such an individualistic focus on hearts and minds is misguided, because it overlooks the central importance of socio-structural factors like housing segregation and family leave policies. However, these discussions implicitly assume that whether something is an individualistic or structural phenomenon turns on whether one can point to an individual mind in explaining it. I argue that this assumption is mistaken. Biased minds are themselves social structural phenomena that arise from the internalization of systematically transmitted cultural information. Misdiagnosing phenomena like implicit bias as individualistic problems overlooks the systematicity that characterizes these phenomena, and so is likely to lead to inadequate prescriptions for remedying social inequality. I therefore propose expanding the scope of what we recognize as structural phenomenon so as to better reason about possible solutions to systemic injustice. More specifically, I offer a dual account of social structure that captures both the external structure that people typically identify, as well as what I call internalized structure. External structure focuses on institutions, and encompasses housing segregation, family leave policies, accessible public transportation, and other factors that structural prioritizers rightly emphasize. Internalized structure, on the other hand, focuses on cultural scripts, schemas, representations and other bits of cultural information that individuals systematically internalize and act upon. By recognizing these internalized structural phenomena as structural, the dual account of social structure that I offer clarifies the nature of problematic phenomena like implicit bias. In doing so, it reveals the inadequacy of individualistic remedies, while also suggesting new avenues for social intervention.

The Social Ontology of Systemic Injustice

Laura Martin—*The Social Ontology of Systemic Injustice*
Columbia University

The idea of systemic injustice is often invoked to describe pervasive, deeply rooted social inequalities, created by unjust social structures. Widespread sexual harassment, for example, has recently been described as a systemic problem; the use of the term systemic in this context indicates that while there are agents responsible for harassment, there are also unjust social structures at work which give rise to this particular form of harm.

Clearly, the concept of systemic injustice is a powerful one for social analysis and critique. Yet its commitments in social ontology are far from unproblematic, particularly with respect to what has been called the problem of structure and agency. On one hand, placing emphasis on the causal force of social structures threatens to result in a diminishment of human agency, and an unacceptable determinism. On the other hand, an emphasis on agency explains neither the patterning characteristic of structural phenomena, nor structural constraint. Whereas common approaches tend to emphasize either the role of individual agents or broader social structures, this paper elaborates and defends a social ontology in which these roles are co-constitutive. I argue that this ontology is necessary to explain how systemic injustice is both perpetuated and resisted by social agents.

One prevalent approach to the structure-agency problem proposes that it can be resolved with the help of rational choice theory. On this view, which can be found in the work of Ann Cudd and Sally Haslanger, for example, unjust social structures makes it rational for those who are disadvantaged to make choices which perpetuate their own oppression. While this approach captures one important dimension of the phenomenon, it places too much emphasis on agents' capacity for rational decision-making. It is therefore ill-equipped to cope with the more habitual mode in which social agents perpetuate oppression. A second family of approaches, found in the work of John Searle, for example, emphasizes the habitual manner in which social agents reproduce social structures. While this strategy avoids the pitfalls of the first approach, it tends to overemphasize the role social structures play in the determination of human action, and hence allots too little room to agents' capacity for reflection, choice, and awareness, even in habitual activity.

In light of these one-sided approaches, I propose a framework in which social structures and social agents are mutually constitutive. At the core of this framework is the following idea: in their habitual, everyday activity, social agents actively draw upon socially embedded rules and resources. In situations in which these rules and resources are unjust (such as the meanings, norms, and power relations common to cases of sexual harassment, for example), social agents can be understood as engaged in reproducing unjust social structures, while retaining their status as agents. This approach creates conceptual space for action which can be seen as shaped by social structures and carried out by genuine agents. It therefore charts a middle path between the two alternatives outlined above, promising a more fruitful avenue for understanding how social agents participate in the creation of systemic forms of injustice.

Populism, Social Groups, and Democratic Health

Eduardo Martinez—Populism, Social Groups, and Democratic Health
University of Michigan

Recent developments in democratic practice and theory hint at cracks in the foundation of support for democratic institutions. Public opinion surveys demonstrate a growing disaffection with democratic institutions and openness to authoritarian government among citizens of established democracies (Foa & Mounk 2016; 2017). In addition, theorists are exploring alternatives to primary democratic institutions, such as elections (Guerrero 2014; Lopez-Guerra 2014), and even democracy itself (Bell 2015). In response to these developments, I articulate an account of democratic health and its epistemic requirements for citizens and the social environment in which they operate. I argue that democratic theorists should adopt this account of democratic health in order to show that democracy can reliably realize its motivating values (and so is not self-defeating), and that it can continue to improve its ability to do so.

I motivate my account of democratic health through an analogy. Much like public health systems require more than a basic infrastructure of care providers and medical suppliers to promote health in a population, democracies require more than electoral institutions and political rights to facilitate self-governance. In particular, democratic citizens must be able to (i) make well-informed political judgments about their interests and how to advance them and (ii) effectively communicate their judgments and relate them to others. However, much like public health officials must address social determinants of health, democratic theorists and practitioners must consider the social determinants of political attitudes and behavior, and the challenges they pose for democratic health. Otherwise, democratic institutions may fail to serve as a vehicle for self-governance, and thereby fail to realize the values that motivate them.

I then argue that at least two features of the social epistemic environment in the U.S. pose challenges to democratic health. The first is populism, understood as a rhetorical strategy that appeals to a social cleavage between underserving elites and “the people.” While such rhetoric is not inherently negative, a social environment characterized by it can obscure heterogeneity among the citizenry, thereby inhibiting citizens’ ability to communicate and relate their interests to one another. The second social epistemic challenge is the difficulty of communicating across lines of social identity. When those with different social identities (e.g., ethnoracial identity) diverge with regard to political attitudes and behavior, the claims made by members of one group can be easily misunderstood and dismissed by those outside the group. For example, public debates about school integration may devolve into misunderstanding and mutual distrust as opponents of integration misunderstand proponents’ use of the term ‘racism’ to describe the status quo (Mendelberg & Oleske 2000). Under these conditions, citizens are prevented from effectively sharing information across identity lines and forming better-informed judgments.

I conclude by demonstrating how an epistemic conception of civic virtue provides the resources to explain what is needed to improve democratic health. This conception captures the interdependence between individual practices and the social environment in which they operate, and how these two features of democracies must align to facilitate self-governance.

Deontic Binding

Russ McBride—Deontic Binding
University of California, Merced

One prominent stream of work within social ontology (e.g., Gilbert, 1990; Searle, 1995, 2010) is committed to what we can call, the Deontological Thesis—the thesis that the most important aspects of social reality ultimately rest on deontology, i.e., rights, duties, obligations, authorizations, permissions, etc. These form the atomic building blocks of social reality. Searle also believes what we can call the Pan-Declarative Thesis—the belief that whenever a deontic power is created, it is created by a Declarative speech act. We now have an interesting argument. If one is committed to both the Deontological Thesis and the Pan-Declarative Thesis then one is necessarily committed to the claim that the most important social reality rests ultimately on the declarative speech act. Unfortunately, there are a wide collection of obvious counter-examples where a social entity comes into existence without a Declarative speech act.

Fortunately, this problem points to an important insight. The source of the problem, I suggest, is that we do not yet have an account of the kinds of deontic structures in play. Searle (2010) claims to have tried and failed. Most of the problems with attempts to formalize a deontic taxonomy stem from trying to categorize the deontics themselves. A distinct approach involves reflecting, instead, upon the ways that deontic powers bind to an agent. There are at least three ways of binding deontic powers to any agent. The first two ways emerge from a distinction between those rights and duties forced upon an agent versus those the agent voluntarily accepts. Within the category of voluntarily accepted deontics, however, there is an interesting subtype which, rather than being created external to the agent, is created, instead, by the agent. There are, then, three important categories of deontic binding: imposed, voluntary, and self-created (autogenic). All the failed cases where there is no Declarative speech act stem from the category of autogenic deontics.

The Intentional Structure of the Modern State

Molly McGrath—*The Intentional Structure of the Modern State*
Assumption College

This paper uses the resources from Searlean social ontology and from Husserlian phenomenology to articulate the intentionality involved in the modern state. Using the tools of intentional analysis from these traditions, this paper examines key texts from the history of political philosophy to argue that there is a distinctly modern concept of state with an interesting intentional structure differing from pre-modern conceptions. In the development of this concept, how are these key terms related: state, government, “the people,” and sovereignty? The term that mediates them is “representation,” but it is a tricky sense of “representation” (related to the modern epistemological sense of the word) of which democratic representation is only one, derivative mode.

The modern concept involves hypothesizing something that cannot be directly experienced—the state—something that can show up only through its representation by people who claim for themselves the ability to speak and act on its behalf, and the ability to make law for it (those who claim to be the “legitimate” government). In order to be successful, for a state in the modern sense to be constituted, these claims to represent must be accepted or recognized by those people to whom the representatives address the law they make.

One motive of the modern theory is to protect the independence of this claim of authority from any other claims of authority. The claim of “sovereignty” is meant to put an end not only to church claims of authority to meddle in secular government, but also to the messy play of parts of the city for control, as we read about in Aristotle’s *Politics*. The key problem of modern political theory becomes to justify and articulate a theory of sovereignty in order to guarantee the governing group’s unique right to make law in the name of the state. Both the concept of the state (as opposed to a political community or realm) and the concept of sovereignty (as self-contained, self-originating, and absolute authority) are essentially modern, with precursors but without proper correlatives in pre-modern political thought.

Insofar as this represented thing, the state, is conceived as existing independently of the representation provided by the government, and accessible only through this representation, it is an occult entity, similar to a thing-in-itself, part of a noumenal realm inaccessible to cross-checking or verification. In terms of contemporary social ontology, it is a covert social construction, one that in order to justify itself covers up its origins as dependent upon acceptance or recognition. One consequence is the systematic ambiguity between “the state” as (purportedly) metaphysically identical to “the people” subject to the law and “the state” as the government, the group claiming uniquely to represent it and thus to rule it.

Texts especially important in this analysis are Searle’s two books on social construction, Edith Stein’s *An Investigation Concerning the State*, Hobbes’ *Leviathan*, Locke’s *Second Treatise*, Rousseau’s *The Social Contract*, and various pre-modern texts (e.g., Aquinas’ “*Treatise on Law*,” Aristotle’s *Politics*) for comparison.

What Determines the Boundaries of a Political Community?

Torsten Menge—What Determines the Boundaries of a Political Community?
Northwestern University

When we think of political communities, we usually have in mind a bounded group of people, people who are related to one another in the right way and who, in most cases, inhabit a particular piece of land. Who constitutes the membership of a community is determined by facts about the members of the community itself (e.g. their attitudes towards one another). A serious problem arises when we apply this understanding of political community to normative questions (e.g. questions about the ethics of immigration): How can we determine who legitimately belongs to a group without begging the question by already assuming who is a legitimate member? But effectively, since this is a difficult problem to address, appeals to political communities in normative debates often treat historical group boundaries as given. This is troubling because it can obscure important normative questions.

Following Brian Epstein's (2015) recent challenge of an anthropocentric social ontology, my goal in this paper is to challenge a political ontology that narrowly focuses on the people that are the members of political communities. I argue that the civic boundaries of political communities are in fact determined not simply by the attitudes and decisions of its members, but also by histories of conquest and colonialism (and their material traces), the material practices establishing and protecting territorial borders, the enforcement of immigration policies, the global infrastructure that enables much of the social and political life in wealthy modern societies, etc. As an example for this point I will focus on the so-called "global land rush," i.e. the recent surge of foreign land acquisitions by governments and corporate investors for the purpose of plantation agriculture and access to water, metals, and minerals (my discussion will draw mostly on work by the sociologist Saskia Sassen (2014)). The global land rush has led to new forms of territorial governance that allow direct control of land and resources by high-income countries, unsettling the usual congruence between the civic and territorial boundaries of a political community. Mediated by global material infrastructures, those who inhabit the newly acquired lands have a significant stake in the political decisions of the countries that now own those lands, raising the question whether those inhabitants are (or at least should have the right to be) legitimate members of those political communities.

I use the issue of foreign land acquisitions to motivate a view of political communities as social-material assemblages, which include not just people and the social relations among them but also the material infrastructures that make much of our social and political life possible. After providing a basic sketch of that view, I argue that some of its counterintuitive implications are outweighed by its ability to make visible important normative issues that are obscured by the "methodological nationalism" of most anthropocentric political ontologies.

The Consequences of the Social Construction of Race for Social Contract Theory

Kristina Meshelski—The Consequences of the Social Construction of Race for Social Contract Theory

California State University, Northridge

In Rawls's original position, the parties to his hypothetical social contract are not supposed to know their class, race, or gender, in order that they will choose principles of justice that are general and unbiased (Rawls, 1971/1999). Feminist and anti-racist philosophers have criticized this aspect of the veil of ignorance for being inadequate to address the reality of gender or racial injustice, and even possibly worsening such injustice by obscuring it (Pateman, 1988; Okin, 1989; Mills, 2005). But so far these criticisms of Rawls have not engaged with contemporary theories of the nature of race or gender, and thus they have unwittingly taken on some of Rawls's essentialist assumptions. In this paper, I will concentrate on the social construction of race, considering many of the most prominent theories and what they would mean for the Rawlsian. (Including biological theories that are not at odds with construction such as Spenser, 2012; Andreasen, 2005; as well as realist constructionist theories like Haslanger, 2000; Alcoff, 2006; and nominalist construction theories such as Glasgow, 2009; Appiah, 1992.) I also note that while critics have concentrated on Rawls, what I say will hold for almost any contemporary social contract theory, to the extent that it requires the deliberation be ahistorical (E.g. Scanlon, 1998; Gauthier, 1986). Though varieties of social construction will have slightly different implications, for any social construction theorist the problem is deeper than earlier critics have noticed. If race is indeed constituted by unjust social history, than the parties to the social contract could not have any concept of race at all, without knowledge of that particular social history. So, rather than simply not knowing their own race, they wouldn't know what it means to have a race. But allowing them to know the particulars of history threatens to collapse the entire social contract project, in so far as such a project is attempting to reveal the objective nature of justice. Interestingly though, the social contract theorist does not have the same problem with class, despite the fact that class is likely socially constructed as well. Considering the differences between the social construction of race and the social construction of class, as they play out for the social contract theorist, may help provide some insight as to why there has traditionally been tension between racial and class politics.

Vagueness and Social Ontology

*Imko Meyenburg and Ana Turcitu—Vagueness and Social Ontology
Anglia Ruskin University*

The aim of the paper is to discuss the relevance of and impact on the philosophical concept of vagueness on the attempt to formulate a coherent social ontology, with focus on the work of Tony Lawson and the Cambridge Social Ontology Group (CSOG). This realist social ontology presupposes that there is a domain of phenomena reasonably demarcated as social reality, or the social realm, which allows studying the nature, the modes of being and the relationships between all phenomena, existents, properties etc. within this realm and whose existences necessarily depends, at least in part, upon human beings and their interactions. Lawson defines the specific process from which social reality emerges 'social positioning', where 'human individuals, things or other phenomena become incorporated as components of these emergent totalities' as a result of commonly acceptance. Moreover, humans are given a special status as their positions operate under interdependent rights and obligations. Finally, under this social ontology it is argued that the systematic study of these elements has significant implications for epistemological and methodological questions in the social sciences.

However, vagueness provides a potential ontological problem. In philosophy, the term is usually associated with the inexistence of clearly defined borderline cases. The classic example is the question of where the heap of sand stops being a heap of sand when individual sand particles are removed one by one. Applied to social ontology, specifically Lawson's social positioning process, vagueness may pose a problem to define borderline cases, for example where questions about the requirements for something or someone to be 'depositioned'. All we are left with are relative borderline cases. In addition to this, and more broadly, vagueness raises the question whether our languages or (social) reality, or both, are fundamentally vague. Both situations will have substantial implication for social ontology by forcing us to accept an unknown number of intermediate states which are inquiry resistant.

Judgment, Agency and Joint Epistemic Action

Seumas Miller—Judgment, Agency and Joint Epistemic Action

Charles Sturt University

For some time Ernest Sosa and others have been propounding what is referred to as performance-based epistemology according to which believing and judging are types of performance with an epistemic aim or end. Thus according to Sosa what he refers to as apt belief is a species of knowledge and in order for a belief to be apt it must be true and true because it manifests the believer's cognitive competence. Here competence consists of three elements: (1) an ability to succeed if (2) the agent in question is in an appropriate inner state and (3) in an appropriate outer situation. Thus the ability to drive is the ability to drive successfully if one is not drunk etc. and the road conditions etc. are suitable for driving. However, Sosa's account, and those of others in this area, are essentially focused on the individual knower or believer; it has not been generalized, nor its implications explored to any serious extent, in relation to collective knowledge. Indeed, given that collective knowledge is standardly understood in static rather than performance terms, this is understandable.

Meanwhile outside the domain of epistemology the notion of joint action has now for some decades been receiving detailed philosophical attention. Consider the numerous different analysis of familiar examples, such as walking together, painting houses together etc. Moreover, the notion of epistemic action has been discussed in the literature including, in a manner, by Sosa. Roughly speaking, epistemic actions are actions directed to an epistemic end, such as knowledge. An example of an epistemic action is a person working out the answer to a crossword puzzle. Recently, these two notions of joint action, on the one hand, and epistemic action, have been brought together to yield the notion of joint epistemic action and analyses have been provided of joint epistemic actions of various kinds. An example of a joint epistemic action is two persons jointly working out the solution to a crossword puzzle. However, these analyses of joint epistemic action have not hitherto benefitted from the extant detailed epistemological work conducted under the banner of performance based epistemology. Given the obvious complementarity between these two strands of philosophical analysis, this is evidently a major omission.

In this paper I undertake three tasks: (1) outline Sosa's account of performance based epistemology with an eye to its incorporation into the analysis joint epistemic action; (2) provide a preliminary analysis of joint epistemic action; (3) elaborate and defend an analysis of joint epistemic actions that incorporates the relevant components and insights of performance-based epistemology.

Intentional Institutional Agency in Oppressive Societies

Anna Moltchanova—Intentional Institutional Agency in Oppressive Societies
Carleton College

I discuss two prominent theories of institutional group agency and offer a view that overcomes their limitations in modeling oppressive societies.

Kirk Ludwig's reductionism would evaluate an oppressive society as not an intentional group agent if the individual engagement in the status roles that engender the patterns of group activity is not, for most persons, intentional directly in relation to the status role they perform. An oppressive group may qualify on Deborah Tollefsen's view as an intentional group agent, because the attitudes regularly ascribed to certain groups are to be identified not with sets of variously interrelated individual attitudes, but with the dispositional states of the whole. Thus, if Canada and North Korea both act according to their long-standing beliefs, they are intentional institutional agents. However, Tollefsen's view, overlooking individual perspectives, wouldn't differentiate between what an oppressive and a democratic country think although they differ greatly in terms of institutional legitimacy.

I argue that oppressive societies are intentional institutional agents, albeit deformed as compared to their avowed organizational principles and defective judged by the standard of what an ideal legitimate institutional group is on Ludwig's view. I outline a view of we-part awareness that is neither an I-awareness nor an undifferentiated "sense of us" and show how using this notion of we-awareness helps us to build social ontology applicable to societies that deviate from the liberal-democratic norm.

I argue that oppressive intentional institutional agents not we-aware with respect to their group functioning. I consider collective self-awareness to be comprised of individual loci of we-awareness. In a large group, individuals are we-aware whenever most of them, from their first-person perspective, experience events as a "we-part" of the we in which they include mostly the same individuals.

In oppressive societies, an intentional institutional group includes both the oppressed and the oppressor. In group actions, the part of the group that forms the intention related to the function of the action in the institutional framework and that controls the action execution has not been delegated these powers by the rest who perform the action and don't intend the outcome. The oppressed individuals participate in joint actions that allow the group to function, but not as part of a "we" under the description corresponding to the officially declared status roles they nominally perform. They can individually coordinate by following instructions or they can act as a we-aware group but in the subsidiary actions, not falling under the official description, that they have to perform to comply with orders.

The oppressed have no we-intention corresponding to the goals of institutional functioning, neither as a "we" of the whole group (including the oppressor) nor as a "we" of the oppressed (excluding the oppressor). Thus, the group is not we-aware but the group acts based on the institutional intentions.

The we-part approach adds first-person experiential awareness to an account like Tollefsen's and shares Ludwig's intuition concerning what oppressive institutional agents aren't, but says more concerning what they are.

What Counts as Fundamental?

*Jonah Nagashima—What Counts as Fundamental?
University of California, Riverside*

One reason to care about metaphysical grounding is its putative connection to fundamentality: if p grounds q , the p is more fundamental than q . If there's a fundamental level to reality, grounding will aid us in discovering what it's like, because grounding captures notions like metaphysical explanation. And fundamentality is often connected with the notion of substantivity; nonfundamental questions aren't metaphysically substantive. But what's on the fundamental level? Here's two standard proposals: some think that the fundamental level contains the constituents of ordinary objects, i.e. the fundamental particles of physics; others think that the fundamental level contains the cosmos as a whole. These are incompatible explanations. So then: What counts as the right kind of explanation? What counts as fundamental?

Notice that agents disappear from the picture either way---either agents are less fundamental than their constituents or agents get subsumed by the cosmos. So, things that depend on agents don't count as fundamental; so, metaphysical questions involving those things don't count as substantive. As a consequence, Elizabeth Barnes and Mari Mikkola have argued that if either proposal is right, then there's an implausible result: much of feminist metaphysics is nonsubstantive. Consider questions like "are there genders?" or issues in the metaphysics of race. According to Sally Haslanger and others, genders depend on social structures, which in turn depend on the activity of agents, and hence aren't part of fundamental reality. So, debates about the metaphysics of gender are (implausibly) metaphysically nonsubstantive (same for debates in the metaphysics of race).

In response, this paper articulates and defends fundamentality pluralism, according to which different domains, corresponding to different (sometimes incommensurate) inquiries, establish different standards of explanation. Roughly, arguments or inquiries in metaphysics establish a set domain such that given the inquiry, there will be objective facts about what facts need to be explained, what facts are eligible to do the explaining, and what the right kind of explanation looks like. Once set, grounding helps to take us to the fundamental level of reality. Hold fixed the inquiry, and there will be an answer as to which view we should favor, but there's no necessary reason to think that one inquiry is more fundamental than the other, so there will be an irreducible plurality of fundamental levels. By bringing in some of the pragmatic considerations involved in explanation, fundamentality pluralism provides a way to account for these kinds of cases while still upholding a robust sense of objectivity about metaphysical truths.

I then provide utility-based arguments for fundamentality pluralism. Centrally, I argue that the view allows us to make space for metaphysically substantive disputes involving things like gender and race. On questions like "what genders are there?" or "what is race?", there's plausibly nothing in Platonic heaven that can provide answers. On many accounts, looking at the subatomic level isn't helpful either—you need to look at societies as a whole. So, societies and social structures can be explanatory stoppers within a domain, which is a mark of the fundamental.

The Division of Ideological Labor

Bernhard Nickel—The Division of Ideological Labor

Harvard University

Ideologies pose the micro-macro problem particularly acutely. An ideology at least contains a set of propositions that concern a group of people and include stereotypes, archetypical narratives, and clichés, but also more general principles that are intended to explain and ground the relatively more surface-level purported facts (Geuss, 1981; Mills, 2005; Shelby, 2003).

Because ideologies are macro-level phenomena that help us understand the persistence of social structures and, depending on the case, unjust distributions of resources and opportunities, they raise the question of how the psychology of individual agents—a micro-phenomenon—is related to an ideology. Ideologies pose this micro-macro problem in a particularly acute way because they can operate without the conscious awareness of the people whose thought and action we understand by invoking them. In such cases, standard pictures of how social-level cognitive elements can be related to individual-level ones, e.g., via the notion of common belief, are inapplicable (cf., e.g. Elder-Vass, 2011; Greenwood, 2003).

This paper proposes a new model for how the psychology of individual agents is related to a sociallevel construct like an ideology, inspired by the paradigm of semantic externalism about reference (Burge, 1991; Kripke, 1972; Putnam, 1975a,b). Because semantic deference generally operates below the level of individual consciousness, it's a promising model for understanding how ideologies function.

The paper goes beyond these more familiar forms of deference in arguing that speakers can defer about the explanatory resources they deploy. The theory of generic sentences due to Nickel (2016) offers a case study of such explanatory deference. It suggests that speakers can be unaware of the ways in which they defer to explanatory external cognitive resources in many different domains. I argue that we can extend this account to address the micro-macro problem for ideologies.

In that case, ordinary speakers who might have only a patchy grasp of the ideology are appropriately characterized in terms of the more worked out ideology produced by what Shelby (2003, 161) calls ideologists, for ordinary agents defer to the cognitive resources made available by these ideologists: public figures who produce op-eds, books, talk-shows, chain emails, etc.

Of course, while in the case of deference about reference, the deference to a common cognitive resource is only a way-station on the road to being in touch with extra-social reality, this is not present in the case of ideologies. But even so, ideologies and ideologists provide the stabilizing and coordinating role that experts do in the more familiar examples.

The crucial contribution of the semantic deference model consists in providing us with an example in 1 which deference to a common resource can be present without the speaker being consciously aware of the fact that she defers in this way, even when it comes to explanatory resources, not just referential ones.

White Supremacy as Metaphysics

Patrick O'Donnell—White Supremacy as Metaphysics

Johns Hopkins University

Despite a few notable outliers (Mills 1997, 1998, Taylor 2004, Curry 2017), there are few serious attempts within philosophy to theorize White Supremacy as a central problem for social ontology, and there is virtually no engagement with the theme within “analytic” social ontology. This paper construes WS as a metaphysical framework within which facts about racial inequalities are grounded. Here “White Supremacy” is conceived not as mere racist ideology, but as the social-structural system of racial domination characteristic of the modern world. Consequently, a suitably theorized notion of WS is supposed to figure into a complete explanation of why racial inequalities persist. How can we make sense of this claim?

The notion that WS can be genuinely explanatory of racial inequalities is often met with skepticism. One skeptical argument is that facts about economics, human psychology, interlocking social practices, history, etc. are sufficient to explain contemporary racial inequalities without invoking any such boogeyman as “White Supremacy” (Lebron 2013, Reed, Jr. 2016). Yet this critique fails to recognize that “WS” is not always presented as a causal or genealogical story about how racial equality came to be and how it sustains itself. Rather, many theorists and activists who use the concept seem committed to the view that WS reproduces social ontological facts by means of non-causal mechanisms, or that it is an underlying framework “within” which various race-specific causal processes play out.

The central move of the paper is to understand WS as a “frame” within which facts about racial inequality are grounded. What sorts of facts does WS ground? First, “WS” talk usually focuses on the fact that Whites and non-Whites are often treated differently within social structures, even when putatively “non-racial” features of the context are fixed (e.g. race-based disparities in sentencing for identical crimes.) Consequently, WS should explain how facts about the normative significance of actions are grounded in facts about racial presentation. Second, “WS” talk often focuses on how differential treatment based on race reinscribes society-wide patterns of racial advantage and disadvantage. Consequently, WS should explain how probabilistic facts about life chances and outcomes are grounded in facts about racial presentation.

I concentrate on conceptualizing the first of these grounding relations, with the understanding that this takes us some way toward understanding the second. Understanding how the normative significance of actions are grounded in racial presentation requires us to understand the role that WS plays in shaping agency. Here my suggestion is that WS can be understood to ground various “frame principles” (Epstein 2015) which specify relationships between racial membership, behavior, and the normative significances of such behavior. Evidence of these framing principles can be found in concrete historical examples-- e.g. restrictive covenants, disenfranchisement, stop-and-frisk-- but these are instantiations of a more general framework in which race/behavior pairs are “mapped” to context-specific normative significances. “Agency” under the conditions of WS is then understood as a kind of “total map” of the constraints leveraged upon non-White individuals and collectives. In closing, I show how this understanding gets us started on thinking about how facts about race ground probabilistic facts about life chances.

The Religious Conception of Social Ontology in Kierkegaard's Concept of Anxiety and Its Implications for the Philosophy of Race

*John Oduor—The Religious Conception of Social Ontology
University of Essex*

Kierkegaard is not a political thinker. However, in this paper I will argue that in *The Concept of Anxiety* Kierkegaard comes very close to developing a social ontology capable of offering an account of the relationship between individual agency and an unjust world. What stops Kierkegaard from developing a fully social, or political, conception of agency is his insistence on the inexplicability of anxiety. What a religious conception of sin and a modern conception of politics share is a recognition of the fact that our individual's agency is deeply implicated in the pervasive wrongness of the world. Kierkegaard's religious interpretation of this implicatedness is understood in terms of our anxious relationship to original sin; in contrast to this, a political conception of agency attempts to understand this implicatedness in socio-historical terms. What unites the religious and political perspective is an attempt to try and make sense of how the individual can understand her relationship to this pervasive wrongness in a way that does not deny her agency, that is, in a way that leaves open the possibility of salvation or social change. However, it is also on this point that the political and religious strategies diverge. Kierkegaard attempts to explain the relationship between the original, world transforming sin, and our subsequent individual sinful acts, through the concept of anxiety; but whilst doing so he insists on the inexplicability of anxiety. In contrast to this, I argue, a properly political conception of agency cannot abide the existence of phenomena that cannot be explained in social terms. After drawing this contrast between what I take to be distinctive about the political and the religious conceptions of agency, I will look at what I take to be the dominant understanding of race, epitomised in the work of the popular American writer Ta-Nehisi Coates and, more recently, in Charles Mills' engagement with critical theory. What I will show in my discussion of Coates and Mills is that the analytic distinction between religious and political perspectives is useful for understanding what Coates' account of racism and Mill's disagreements with traditional critical theory share. I will argue that the aforementioned author's insistence on the irreducibility of race to other social phenomena such as class is akin to the Kierkegaardian insistence on the inexplicability of anxiety. To this extent their thought, I argue, ought to be understood as religious rather than political. I then conclude by making some tentative suggestions about what place, if not in politics, Coates' and Mill's account of race ought to have in our attempts to understand our relationship to the social world.

Interpreting Collectively Accepted Social Norms from a Perspective of Performativity

Noriaki Okamoto—*Interpreting Collectively Accepted Social Norms from a Perspective of Performativity*
Rikkyo University

Several scholars have investigated in depth the nature of objective institutional reality and suggested the possibility of its application to the social sciences (e.g., Searle, 1995, 2005, 2009; Tuomela, 2002, 2013). For example, the simple but persuasive accounts of Searle (1995, 2005, 2009) have attracted great interest from social scientists and philosophers alike. His theory of social and institutional reality is useful in explaining epistemologically objective reality, but relatively few social scientists have thoroughly applied his core concept of collective intentionality to their specific domains (for example, De Soto, 2001). One of the reasons why Searlean collective intentionality and associated perspectives have not been extensively used in social scientific research is that the concept is considered too abstract and simple to explain the complicated dynamics of social organizations and institutions. Perhaps his main concern lay in ontological clarification rather than extensive application of the term to various social issues, meaning no systematic theory of collective intentionality was developed that could be used in social scientific studies. Another leading scholar of collective intentionality, Raimo Tuomela, presented his latest theory of social ontology on the basis of his perspective of the we-mode and collective acceptance (Tuomela, 2013). One recently published book (Preyer and Peter, 2017) consists of a critical discussion of Tuomela's theories, mainly grappling with the explication of Tuomela's we-mode perspective. As a social scientist, my attention is directly focused on the assertion that "group agents have intentional attitudes only in virtue of the members' reflective performative collective construction of contents for the group" (Tuomela, 2013, p. 15f.).

In social science fields such as economics and sociology, several social scientists have recently directed their attention to the concept of performativity. Several studies in Boldyrev and Svetlova (2016) have focused on the origins of performativity with reference to the perspectives of Austin and Searle. This study integrates some advances in the theory of institutional reality, backed by the concept of collective intentionality, with insights emerging from performativity studies in the social sciences. With a particular focus on the extant literature of collective intentionality from a perspective of performativity, the study attempts to shed light on one aspect of the performative content of collective intentionality. Specifically, Searle and Tuomela's collective acceptance view is reconsidered from a perspective of performative perlocutionary forces in Austin (1962) and Svetlova (2016) in order to suggest a kind of meso-level collective social norm that performatively mobilizes a change in institutional reality. The suggestion assumes that collective intentionality involves normative forces that drive individual and organizational behaviors. It cannot be denied that collective intentionality could have normative deontological aspects, and interesting work has sought to identify such a normative mechanism, such as Tuomela's classification of r-norm and s-norm characteristics (Tuomela, 2002), but both philosophers and social scientists have tended to avoid identifying collective intentionality's consequences and its positive and negative influences. Arguably, the concept of collective intentionality points to the relativity of phenomena or concepts that are believed to be quite solid and objective. Searle (2009) eloquently calls this aspect its "observer-relativity." To identify such "observer-relativity," this study argues for the adoption of a viewpoint of performative collective norms in social sciences that entails the

observer's interpretation of collective norms. In this study, as one of the representations of such performative collective norms, the societal expression of "-zation" (such as globalization, financialization, and marketization) works performatively in society. Analyzing and interpreting the reflexive mechanism of collective norms as such "-zations," this study relies on the logic of perlocutionary forces as per Austin and Searle. Finally, the study touches on MacKenzie's (2006) counter-performativity to explain the dynamics of social institutions.

The analysis from this specific social scientific standpoint can at least be useful to related studies that have considered and analyzed the fundamental essences of collective intentionality, norms, responsibilities, and so on.

"Some Are Guilty, All Are Responsible"

Alex Ozar—*"Some Are Guilty, All Are Responsible"*

Yale University

Taking a cue from Abraham Joshua Heschel's aphorism that with respect to the ills of liberal societies, "some are guilty, all are responsible," and building on the work of Iris Marion Young and others, I develop an account of responsibility for individual agents under conditions of collective structural injustice.

Responsibility is often understood in terms of liability – to say that agent X is responsible for action or state-of-affairs Y is to say that Y was in some way within the sphere of X's agential control – but it can also refer to a prospective relation to objects of value, as in the responsibility I have for my son. I argue that this form of responsibility cannot be reduced to any set of discrete duties, but is rather a fundamentally open-ended, holistic call to seek the object's good. For this reason responsibility is not exhausted by the absence of discrete fault.

While responsibility is for objects of value, I argue that it is to the community of valuers. Since valuing X entails valuing X's being valued, we can hold each other and ourselves to account for any shortcomings in our valuing of X, on pain of exile from the community of X-valuers. Since valuing X is inconsistent with indifference to seeking the good of X, we are accountable not only for our thoughts and feelings but for our actions in seeking that good. Since we are not only impersonally obligated but also (inter)personally accountable, we have a positive responsibility to make our responsibility for X publicly manifest. To take responsibility for X, then, is to actively seek X's good in a manner demonstrably recognizable as such.

Whether we succeed is often ambiguous. We may seek X's good and yet fail as a matter of the public record, as in a child-valuing lorry driver who strikes a child. Our actions advancing X's good may be interpretable as merely self-promoting. To fully meet our responsibility for X, therefore, requires the assumption of cost – sacrifice or risk thereof – in seeking X's good to a degree inconsistent with not valuing X.

With respect to structural injustices for which we as individuals are not at fault, we have a responsibility to honor the value of those suffering that injustice in seeking their good. We take this responsibility, in the first place, by demonstrating our concern through the assumption of cost in taking actions, however preliminary, toward seeking their good. In this way we not only signal our virtue but witness to our concerned recognition of an interpersonally objective value, which is to witness to the reality of that value and so to invite others to join us in our concern. Were enough others to join us, injustices could be effectively meliorated, and so our witnessing is indeed seeking the good of – taking responsibility for – those suffering injustice.

Finally, since such witnessing is necessary for addressing structural injustice, we have a standing responsibility to ensure our public discourse not be such as to render such witnessing unrecognizable as such.

To What Extent Is a Group an Individual?

Rohit Parikh—*To What Extent Is a Group an Individual?*

City University of New York

Dennett in his *Darwin's Dangerous Idea* (1995) and *Kinds of Minds* (1996) discusses an evolutionary hierarchy of intellectual progress. He calls the hierarchy the 'Tower of Generate-and-Test,' where there are five kinds of creatures. These are:

- (1) 'Darwinian creatures,' organisms which are blindly generated and field-tested,
- (2) 'Skinnerian creatures,' which try different responses until one response is selected by reinforcement.
- (3) 'Popperian creatures,' who can preselect an action from many options before doing it in the outer environment.
- (4) 'Gregorian creatures' who import mind tools from the outer cultural environment to construct better inner environments. And finally:
- (5) Creatures like human beings who use these mind tools and, even more, 'language,'

One could ask, "at what level, if any, do groups belong" if indeed we can regard them as individuals or as intentional beings? Since they do use language, one would think, they are creatures of level (5). But difficulties arise in thinking of groups as even Popperian.

In order for a group to have a real identity, it needs coherence in its "views" and in its actions. To think of it as a game theoretic opponent (or partner) one needs a certain amount of predictability. Questions like "What does this group think?" or "What does it want?" need to be addressed.

Such predictability is not always absent. We know quite well how "Russia," thought of as an agent, will respond in case of a nuclear attack. But "the Republican party" or "Republicans in Congress" might be less predictable in their response to say the election of Conor Lamb.

Two kinds of theoretical issues thus arise. One is epistemic coherence which can exist only if the group possesses mechanisms for intra-group communication. An army preparing to wage a battle needs scouts to gather information and to transmit it to troops. A university needs an internal email system.

The other is exhibiting coherence of views where issues like the Arrow theorem or the judgment aggregation paradox may arise. A group where power is more concentrated, at the extreme in one individual, is likely to be more predictable and more consistent in its response. This is the sad message of the Arrow theorem and applies to China or to a religious group with a strong leader. It applies less to a more diverse group like the Democratic party. There is also the issue of consistency over time – something we will not address in this first pass.

So a better question to ask than "do groups exist?" is "to what extent is a given set of individuals (with a name) an individual, and on what issues?" In other words we suggest an algorithmic and game theoretic alternative to the ontological question. We will offer some answers while avoiding a surfeit of mathematics.

Relevant references include Kenneth Arrow, on his theorem, the judgment aggregation paradox of List and Pettit, *Infostorms* by Hendricks and Hansen, and Parikh's work on Social Software.

In Defence of a Metaphysics of Race

Clarisse Paron—*In Defence of a Metaphysics of Race*

University of Alberta

In the philosophical literature on race, a central debate is whether we should define a metaphysics of race. I begin my argument with an overview of the metaphysical debate on race and the four main philosophical positions: a metaphysics of race where, "race is biologically real, socially real, or not real at all"; and a deflationist position that argues against ontological discussions of race (Hochman 2711). In this paper, I argue for the sociopolitical necessity of having a metaphysical account of race because it allows us to define race for not only our theoretical purposes, but so we can reshape our concept of race to accomplish moral and political goals—that is, not just asking "What is race?" but also "What do we want race to be?" (Haslanger, *Gender & Race* 52). In section I, I refute Mallon's deflationist critique that metaphysical accounts of race are apolitical, non-normative, and purely semantic. Mallon argues that the different metaphysical positions of race fundamentally agree on a broad base of metaphysical facts—what he calls the "Ontological Consensus". Although these philosophers agree on many facts about race, the metaphysical debate is semantically sustained since each position argues using a different theory of reference (causal or referential). Mallon's deflationism attempts to redirect the metaphysical debate from discussion of racial ontology to, instead, "determine whether or not we should try to continue or end race-talk, ...race-thoughts, and racialized structures and practices" (Hochman 2725). While I agree with Mallon that race theorists ought to focus their theories around moral and political issues, I ask: how can we address the normative concerns of race-talk and racism without acknowledging the reality of race? I argue that a metaphysics of race is necessary to understand how "race" is really conceptualized and how critical a concept of race is if we wish to understand and reconceptualize race-talk.

In this paper, I specifically align with a social constructionist perspective of racial ontology because it supports a conception of race that aligns with normative uses of "race" and "race-talk", but, more importantly, it underscores how race and racism are real. In section II, I draw on Kaplan and Potter's arguments to illustrate how social kinds can explain real, predictable, biological causal patterns; thus, reinforcing the plausibility of a social constructivist racial ontology. Their examples highlight the social disparities between Black and White Americans are not due to "biological racial essences" but are caused as a result of social concepts of race and racial discrimination. In the final section, I use Sveinsdóttir's conferralist account of sex to understand how certain social, value-laden qualities might be "conferred" onto objective properties of race and racial concepts. This understanding of race, I argue, helps account for different objective properties that correlate to different races and racial identities—yet, also explain the ways that certain races and racial concepts are socially valued (or disvalued) beyond these objective properties.

A Normative Hylomorphic Theory of Social Objects

Asya Passinsky—A Normative Hylomorphic Theory of Social Objects
Dartmouth College

It has been noted in the social ontology literature that the creation of social reality is puzzling, but it is not entirely clear what the puzzle is supposed to be. In this paper I articulate in clear and precise terms a metaphysical puzzle concerning the creation of social objects such as borders and states. After examining and rejecting some solutions to the puzzle that draw on existing views in the metaphysics literature, I develop a novel theory of social objects that provides the basis for a more adequate solution.

The puzzle can be articulated by focusing on a simple imaginary case. Suppose that two people living on a desert island agree that River is to be the border between their respective territories. They thereby bring into existence Border. Each of the following premises is independently plausible, but they are jointly inconsistent:

- (1) Border is a concrete object in the external world.
- (2) Border is brought into existence by agreement.
- (3) Agreement cannot bring into existence concrete objects in the external world.

Premise (1) is supported by our intuitive judgments about the spatiotemporal and causal properties of Border. Premise (2) is supported by our intuitive judgment as to what brings Border into existence under the envisioned circumstances. And Premise (3) is supported by a general metaphysical principle that says, “Thought and talk alone cannot bring into existence concrete objects in the external world.” Since an act of agreement consists of nothing more than thought and talk, it follows from this principle that agreement by itself cannot bring into existence concrete objects in the external world. But this is inconsistent with Premises (1) and (2), which entail that in the envisioned scenario, a concrete object in the external world is brought into existence by agreement.

I consider four possible solutions to the puzzle that draw on existing views in the metaphysics literature. The first is an immaterialist solution, which denies Premise (1) on the grounds that Border is an abstract object. The second is an eliminativist solution, which denies Premise (2) on the grounds that Border does not exist. The third is a materialist solution, which denies Premise (2) on the grounds that Border is identical to River, and since River is not brought into existence by agreement neither is Border. And the fourth is a plenitudinous solution, which denies Premise (3) on the grounds that concrete objects in the external world come into existence all the time, including when we make agreements. I argue against each of these solutions in turn.

The theory that I go on to develop is a neo-Aristotelian one on which social objects are hylomorphic compounds whose distinguishing characteristic is that their form is normative. Border, in particular, may be understood as a hylomorphic compound whose matter is River and whose form is a certain normative role. This theory provides the basis for a solution to the puzzle that denies Premise (3) on the grounds that certain concrete objects (*viz.* hylomorphic compounds whose matter is concrete and whose form is normative) can be brought into existence through the

exercise of normative powers. At the same time, it concedes that the exercise of such powers cannot create such objects *ab novo*, thereby accounting for the intuitive appeal of Premise (3).

Speaking for Another

Grace Paterson—*Speaking for Another*
Stanford University

This paper is concerned with a kind of speaking for another which is exemplified in proxy speech (eg. proxy consent, proxy voting). Here, the speech of one agent in some sense counts as a speech act of another agent such that the agent spoken for (henceforth the principal) is subject to various obligations and entitlements as if they had themselves performed the illocutionary act in question. Consideration of different cases reveals that the way in which the principal is bound by a speech act performed by proxy can vary. For instance, in some cases it may be considered rational for a principal to attempt to sabotage, undermine, disavow, or otherwise distance themselves from a speech act that has been performed for them. These sorts of distancing acts may sometimes be considered rational and perhaps called for even when the proxy was properly authorized, and even when the principal is nonetheless considered by all involved to be responsible for fulfilling the obligations constitutive of the illocution.

These observations suggest that in some cases of speaking for the principal is a proper agent of the speech act (ie. the speaker), while in other cases she is somehow the subject of it in a less demanding respect. To understand how this may be, I distinguish between two sources of normativity which are at play in speech acts. The first is derived from our social and linguistic practices and consists largely in the various commitments and responsibilities one takes on in virtue of a speech act's illocutionary force. The second derives from the fact that illocutionary acts are subject to the same norms of practical rationality that bind all actions. I argue that it is possible to bind an agent in the social-linguistic sense without thereby binding them in the practical sense.

In order for a proxy to bind a principal in a social-linguistic sense, all that is required is that they be authorized to speak for the principal in virtue of the social practices in which they and the principal (and the speech act) participate. Social structures may grant this sort of authority in a variety of ways, not all of which are responsive to the principal herself in any substantial manner. In order for a proxy to bind a principal in the practical sense as well as the practice internal sense, however, they must be both authorized to do so in a way endorsed by the principal, and also committed to making the actions performed conform with the intentions and reasons of the principal just as the principal herself would. When these conditions are properly met, the speech act functions for the principal like any other action she performs and is therefore constraining for her with respect to all the usual practical norms.

The Nature of the Political Association

Armando Perez-Gea—*The Nature of the Political Association*
Yale University

The nature of the political association (commonly called the “state”, but also the “community” or “society”) as a distinctively unique kind of social group needs defending. We live in a world with multinational trade agreements, international defense pacts, international non-governmental organizations, world-wide communication networks, and global corporations. With all these entities that transcend political boundaries, it might seem hard to justify the existence of the seemingly antiquated political associations that populate the Earth. Political associations have been prevalent throughout human history; city-states, kingdoms, republics, empires, and nation-states have been staples of human history. But the fact that political associations were useful in the past does not guarantee that they still serve a purpose today. This critique against political association is not merely proposing an alternative form of political association (namely, a world government), it is saying that political association has run its course. It doesn’t find anything unique of political association that is worth preserving.

While the defense of the political association might feel like a Quixotic quest embarked by an idiosyncratic individual (an accusation that might not be necessarily false), there are some unsuspecting scholars who are my fellow adventurers. For starters, the scholars who defend Shapiro’s planning theory of law, particularly as presented in *Legality* (2010) and “The Planning Theory of Law” (2017), are committed to the idea of social planning as a distinctive form of planning. A key feature that distinguishes social planning from other forms of planning is that social planning is tied to a community (what I call a political association). This relation to the political association is required for a plan (or related phenomenon) to be a law. If the political association is not distinct enough from other social groups, then there is the real risk of an over-expansive definition of political association. The more social groups that are counterintuitively classified as political associations, the more cases of counterintuitive classification of planning as law. We might end with the planning of the cooking club actually being law, a conclusion that would be deeply uncomfortable for a reader of *Legality*.

The particular claim that this paper explores is the origins of the political association from other social groups. The origin of the political association reveals important aspects of the nature of the political association. This argument of origin is a political association’s construction profile. Of the four parts of the construction profile presented by Epstein in “What are Social Groups? Their Metaphysics and how to Classify Them” (2017), an argument of origin deals primordially with the “coming to exist in a world” but also delves into the other three parts, in particular the “continuing to exist in a world”. My position is that a political association maintains the public realm that provides for objects that are constitutive of human flourishing. This is achieved by building on other social groups, in particular a self-sufficient association. This genesis not only keeps the underlying social groups, but also provides foundations for other social groups to be formed.

Visual Perception, Social Institutions and Corruption

*Laura Perez—Visual Perception, Social Institutions and Corruption
Cornell University*

It is plausible to consider that states and events such as beliefs, intentions, desires and perceptions constitute the social environment, in particular, social institutions. Such states and events would partly determine the kind of entity that social institutions such as university and government are. For instance, the intentions of university members that the university acquire, transmit and disseminate knowledge is a constitutive part of its existence. For its part, citizens' beliefs that the government of their country regulates and coordinates economic systems, educational institutions, the police, the militia, and so on, is a constitutive part of government existence.

A route to analyze social institutions such as university and government is in terms of organizations with goals and particular processes which lead to fulfilling such goals. Processes are realized by a structure of roles embodied by human beings. Particular processes which allow the fulfillment of the institution's goals, nevertheless, can be corrupted. For instance, we would say that the university has been corrupted when processes that fulfill the end of acquisition of knowledge are determined fundamentally by market interests. For its part, we would say that government has been corrupted when abuse of power permeates processes that fulfill the end of regulation and coordination of police service in a systematic way.

The general question that motivates this inquiry is: What are the ways in which events and states such as people's beliefs, intentions, desires and perceptions constitute entities such as universities and governments when their processes have been corrupted? In this piece I focus on the kind of constitution that events of visual perception provide to institutions such as university and government, whose processes have been corrupted. To do so, I draw on photographic work and images representing the theme of institutional corruption in university and government settings.

The piece is divided into three sections. In Section 1 I elaborate on the relation of constitution between the kind of entity that social institutions such as university and government are, and events of visual perception. In Section 2 I analyze events of visual perception of situations involving institutional corruption in university and government settings. Finally, in Section 3, I study photographic work and series of images on institutional corruption in university and government aiming to address the question on the relation of constitution between social institutions and events of visual perception. To do so, I draw on images of university brands and work by photographers Daniela Rossell, Yvonne Venegas and Alejandro Cartagena on the privileged lives of political elites and landscapes related to bureaucracy in government.

We-Intentions and Immunity to Error Through Self-Misidentification

*Björn Petersson—We-Intentions and Immunity to Error Through Self-Misidentification
Lund University*

When I intend to raise my arm, your lifting it does not fulfil my intention. The intention is only successful if **I** raise my arm by way of **this intention**. That does not imply much about the nature of my intention to raise my arm. For instance, it is still an open question whether ‘I’ and ‘this intention’ must figure in the content of my intention, as e.g. Searle (1983, 11-12) thinks, or if the self- and agent-referentiality has more to do with the perspective from which an action intention is held, as Recanati (2007, 129) and others argue.

An objection against placing the reference to me in the content of my intention is the argument from immunity to error through self-misidentification. While I may be mistaken about what I intend to raise – experimenters may have tricked me into intending to raise the rubber limb in front of me in the belief that it is my arm, for instance – it seems absurd to ask me how I know that I am the one intending to raise my arm. I can be in error about what I intend but not about who intends it, or so the argument goes. The impossibility of misrepresentation may indicate that there is no representation to begin with – the ‘I’ was never represented in the content of my intention. This is in line with the more general conclusions Wittgenstein, Anscombe and others have drawn from this kind of observation.

This makes things complicated for accounts that place a necessary reference to the agent in the content of we-intentions, like Bratman’s or Ludwig’s. If I-intentions by their nature do not need such a reference in their content, we-intentions would seem to belong to a conceptually distinct kind of intentional states on those accounts. This goes against e.g. Ludwig’s (2016, 189) denouncement of the idea that an account of individual and collective agency requires two distinct concepts of intention, and it undermines Bratman’s (2014) claim about parsimony.

On the other hand, when a we-intention like ‘we intend to lift the boat’ occurs in my mind, I can be mistaken about exactly who intend to lift the boat. I may misrepresent my group, and that seems to prove that the intentional state contains a representation of the collective agent. That would be an argument **for** a ‘content’-account of collective action intentions.

Some varieties of self-misidentification are possible even in the individual case (like misidentifying parts of one’s body, or quasi-remembering). I argue that the ways in which we-intentions are open to self-misidentification resemble some of those cases. My tentative conclusion is that there is one sense in which we-intentions **are** immune to error through self-misidentification. In the individual as well as the collective case, I argue, the most reasonable explanation of the immunity is that the necessary reference to the agent that must figure in a description of an intention’s satisfaction conditions follows from a perspectival feature of the intention, rather than from an element in its content.

Responsibility and Freedom in Community Activism: A Case for Sustainable Development in Vietnam's Emerging Civil Society

*Lien Pham—A Case for Sustainable Development in Vietnam's Emerging Civil Society
University of Technology, Sydney*

Drawing on a research project about community activism in Vietnam, this paper analyses the notion of responsibility in enabling social change. Interview data with 48 Vietnamese community workers suggest that the Vietnamese habitus imposed on these activists a set of moral obligations to families, communities and the State, which shaped their conception of freedom. The majority of interviewees valued community work because helping others particularly those within the immediate families and communities are laudatory and obligatory in Vietnam. On the other hand, the State's pervasive presence in all aspects of social lives enhanced the high Confucius master/servant role, which resulted in limited desire for political participation at the grassroots level. Freedom was viewed by these Vietnamese activists as being derived from responsibility – prospective responsibility to satisfy their obligations to their families, communities and the State before they could move to satisfy their individual rights and freedom. In this way, responsibility became part of the person - their self-consciousness - and thus not seen as a constraint to freedom since they voluntarily imposed on themselves the responsibility before they exercised their freedom (Ballet et al. 2013). At the same time, these activists also had rights and freedom within their social contexts, which informed their actions in their obligations to others, although this was more visible in foreign non-government organisations. This conception of retrospective responsibility - to be accounted for by a person when their freedom has been exercised - takes the individual as a source of moral decisions based on the view that objective freedom exists outside of the person (Sen 1999). These activists' political ambitions and will to exercise agency depended on their conception of morality and freedom that mediated between these two forms of prospective and retrospective responsibilities.

Along with Ricoeur's (1992) dialectic of the self and other, this paper argues that thinking about collective responsibility in terms of civic actions should include prospective responsibility in specific social structures and relations via a person's commitment for others, in addition to retrospective responsibility as accountability of their individual actions. The reasons are threefold. First, prospective responsibility recognises the social structures and relations surrounding a person, and that their freedom and thus power for actions are enacted upon the values which are conditioned by these social factors. An exercise of such power, which may or may not happen, can lead to collective actions because it results from their capacity to be a responsible agent "within their world". Second, sustainability requires not only collective actions, but responsibility that embraces consideration for others' wellbeing now and the future. This implies capacity of agents to impute responsibility upon themselves for others ahead of their own interests. Third, the idea of self-consciousness is worthwhile for understanding oneself as an integral part of society and Nature as natural order of the world. Social development policies that engage with the idea of a person's life as one of Nature can procure sustainable initiatives that fit in with the people's way of life and balance private and public interest needs.

Universal and Particular Standards in Negligence Law: A Journey into the Epistemology of the Act of Judgement

*Orlando Poblete—A Journey into the Epistemology of the Act of Judgement
Feliu y Asociados*

This paper will examine the case for establishing a personalized standard in negligence law from the ontological nature of the act of judgement. I will argue that the act of judgement –essential to adjudication or dispute resolution- is a dispositional quality that is built up by social interaction. My aim will be to show that a personalized standard theory, as that proposed by Ben-Shahar and Porat, fails to grasp the true nature of the act of adjudication: the act of adjudication does not consist in a solipsistic propositional act, but rather in a socially nurtured dispositional knowledge.

Beyond Pluralism: Racial Concepts, Pejorative Content, and Racial Eliminitivism

Darien Pollock—Beyond Pluralism: Racial Concepts, Pejorative Content, and Racial Eliminitivism
Harvard University

In the past decade, a number of philosophers of race have forwarded different versions of what has come to be called “metaphysical racial pluralism” (racial pluralism). Roughly, this view holds that racial kinds exist in the world, if construed in terms of non-pejorative (non-normative) semantic content (Hardimon, 2017; Mallon, 2016; Spencer, 2014).

Racial pluralists are responding to a prominent view in the metaphysics of race literature referred to as “racial eliminitivism” (Appiah, 1998; Glasgow, 2009). Racial eliminitivists argue that racial kinds do not exist; and, as a consequence, racial concepts are empty (i.e., these concepts do not pick out items that exist in the world). Racial pluralists respond to this general position by not only arguing ways in which racial concepts legitimately refer but also ways in which these concepts can be “engineered” to be void of innocuous, normative content.

In this paper, I forward three claims against the racial pluralist view.

First, I argue that engineering racial concepts (even for scientific purposes) does not guarantee that pejorative (normative) content will not affect reasoning about putative racial groups. To support this claim, I consider evidence from social psychology about implicit racial bias (Banaji & Greenwald, 2013) and recent work in philosophy of mind and perception about racial attitudes (Siegel, 2017). The core of my argument is that we have good empirical reason to believe that there is a strong psychological association between descriptions of racial concepts and mental states that have normative import (e.g., beliefs, attitudes, patchy endorsements).

Second, I argue that, given that racial pluralists cannot segregate racial concepts from pejorative (normative) content, the only kind of racial groups that pluralist racial concepts pick out are non-biological racialist racial groups. I defend this reasoning by appealing to a conceptual distinction that racial pluralist Michael Hardimon highlights in his work. Hardimon argues that the ontological extension of what he coins as the racialist race concept is empty not only because it fails to pick out genuine biological properties; but, more importantly, because its semantic content is inconsistent with “the principles and findings of contemporary biology.”

Finally, I establish that, since racialist racial groups do not exist, races do not exist. This final claim is simply a fine-grained revival of the classical racial eliminitivist position. In other words, if it is not possible for the biological semantic content to be segregated from the pejorative (normative) semantic content of racial concepts, these concepts do not refer to any human groups that exist; therefore, racial eliminitivism is true.

How We Act Together

Matthew Rachar—*How We Act Together*
City University of New York

Previous empirical research (****) suggests that common intuitions about acting together in a strong sense involve three judgments: (i) participants in a collective action may lack individual participatory intentions, (ii) exiting a collective action involves an obligation to make that exit public, and (iii) this obligation is present in “morally wrong” cases of collective action. This presentation sketches a theory of collective intention and action that accounts for these intuitions.

Building from Velleman’s view of intention as a representation with a particular content and causal role (1997), I claim that a collective intention is a public representation that commits participants to a course of action. The key ideas are that non-mental representations that play the appropriate roles and have the right contents are intentions and that an expression of a conditional commitment by each participant can combine into a single, categorical commitment. After filling out this claim by clarifying the idea of a public representation, specifying how public representations relate to the idea of collective commitment, and spelling out how these representations fulfill the action-guiding roles of intentions, I show how this account incorporates and explains the three guiding judgments. First, public representations are created by individual representational acts and persist until they are rescinded or fulfilled, rather than being realized in the minds of the participants throughout the collective action. As a result, they act indirectly on individuals; they are mediated by an individual’s beliefs or memories about the initial creation of the collective intention. A participant can then be party to a collective intention without having a participatory intention, which explains (i). Second, leaving a collective commitment is more difficult than rescinding an individual commitment because it also involves a representational act that indicates that the condition on the other participants’ conditional commitments is not satisfied, which explains (ii). Third, the utterance of a conditional commitment in the presence of other conditional commitments creates the obligation to fulfill or rescind the commitment, regardless of the act in question, which explains (iii).

A Puzzle for Social Essences

Mike Raven—*A Puzzle for Social Essences*

University of Victoria & University of Washington

This paper raises and explores a puzzle concerning social essences.

The social world contains many kinds of entities. Among them are institutions (nations, clubs), groups (races, genders), and objects (ceremonial daggers, borders). It was once common to investigate social entities by investigating their essences or natures. On this approach, discovering what a nation, or race, or border, is involves discovering the essence of a nation, or race, or border. The approach, however, was often combined with problematic presuppositions about what these essences had to be. It was once presupposed, for example, that the essences of races had to be biological. This presupposition turned out to be scientifically unsupported, philosophically unsubstantiated, and socially unjust. But it is increasingly recognized that such presuppositions are not essential to the approach itself. And so some have recently attempted to revive the approach without them.

I wish to explore a puzzle that arises for social essences even after the removal of these problematic presuppositions. The point of departure concerns where essence “resides”. Two conflicting approaches have endured since antiquity. The first (“Platonic”) approach regards essences as somehow detached from the world. It’s as if essences are given prior to their worldly manifestations, as with Plato’s Forms. The second (“Aristotelian”) approach regards essences as somehow embedded in the world. Essences are as much a part of the world as their worldly manifestations, as with Aristotle’s universals.

The general conflict between the approaches is starkest in the case of social entities. In the Aristotelian spirit, not only do we construct social entities, we also construct their natures. Clubs, races, and borders did not predate us. We made them. And we made them what they are. It is, for example, essential to a border that it demarcates a nation’s land. It is so because, somehow, we made it so. The essential fact about borders seems embedded in worldly facts about our worldly institutions, needs, interests, and practices. But in the Platonic spirit, the essences of things do not seem up to us. Essences are like definitions or axioms. Even if we might adopt different definitions or axioms for different purposes, it is not as if we can change what the definitions or axioms are. Nor can we determine what the very natures of things are. Essences seem detached from any worldly grounds. The puzzle, then, is that we are at once inclined to regard social essences as detached from but embedded in the world.

My first aim is to clarify what the puzzle is. This involves refining the notions of detachment and embedment used to state it. My second aim is to explore some of the puzzle’s implications. The puzzle, once refined, turns out to be surprisingly resilient. Engaging with it promises to be of interest for what it might teach us about social essences in particular. It also promises to be of broader interest as a case study of the general conflict between the two approaches to essences.

Bayesian Trust, Trust Intermediation & Boundary Spanning

*Rene Reich-Graefe—Bayesian Trust, Trust Intermediation & Boundary Spanning
Western New England University School of Law*

Williamson's economic model of trust (Williamson, O. E. [1993]. *Calculativeness, trust, and economic organization. The Journal of Law and Economics*, 36(2), 453-486) makes the claim that the term "trust" is regularly misused in social sciences research for behavioral phenomena which are better explained in terms of calculativeness. Prior contributions have modeled the inescapably calculative and institutionalized nature of trust and have explained how individual and collectivized trust expectations and preferences are 'bullet-proofed' into the self-referential and autopoietic process of trust transactions. They also explained how humans are conditioned to calculatively trust for purposes of (economic) cooperation — a process which crucially requires an organized masking of the calculative nature of their trust transactions — and claimed that trust may therefore be best understood as double-blind calculativeness. As a result, and given the carefully shielded, (self-)deceptively calculative nature of trust relations, it is not calculative trust, but its exact opposite — non-calculative trust — which is truly oxymoronic, i.e., "a contradiction in terms" (Williamson, 1993, p. 463).

The current paper further describes and tests such novel economic conceptualization and explanation of all trust transactions as double-blindly calculative behavior. It develops an integrated model of Bayesian trust, trust intermediation, and trust boundary spanning that organizes and explains the multidimensional intra- and interpersonal layers of trusting behavior and describes the calculative modes for the optimization and overall predictability of such behavior. Doing so, the paper also identifies and isolates the microfoundations of personal trusting behavior which then become the collective building blocks of intra- and interorganizational trust relations and human cooperation, thereby enabling the ultimate creation and replenishment of institutionalized, social trust as a public good. In effect, the model radicalizes established social-sciences conceptualizations of trust and, through applied examples of Bayesian updating of trust beliefs, abandons the conventional separation and differentiation of trusting behavior from calculative, opportunistic modes of cooperation in mainstream economic theory.

Saying What Race and Gender Are

Jared Riggs—*Saying What Race and Gender Are*
University of Toronto

In this paper, I defend what I call coordinating expressivism about the meaning of claims like “there are more than two genders” or “there are no races.” Coordinating expressivism is the view that the meaning of such claims is an attitude that 1) need not be fixed by semantic rules of the language prior to any particular conversation and 2) need not be propositional. In some conversations, to say that there are more than two genders may be to express the belief that more than two things realize some theoretical construct, viz. gender. In other conversations, to say that there are more than two genders may be to express a plan not to reject people’s claims that they are neither a man nor a woman.

I argue that coordinating expressivism is superior to extant understandings of claims of this sort. For instance, we might endorse a purely descriptivist interpretation of such utterances, so that a claim like “there are no races” means something like: nothing meets the semantic criteria for the term ‘race.’ But this view has a hard time vindicating the view that people who associate different criteria with the term ‘race’ can have legitimate disagreements over whether a sentence like “there are no races” is true. Such disputes would have to be purely verbal, turning on a disagreement over what the term ‘race’ means. But this is an unpalatable view to hold.

Alternatively, one could understand claims like these as metalinguistic negotiations, where speakers employ semantic content fixed by prior rules of the language in order to express disagreements about what meaning ought to play the cognitive role of concepts like black or woman. On this view, when two people disagree over how many genders there are, they can be interpreted as either disagreeing about some pre-established semantic content or disagreeing about what meaning ought to play the role of the gender concept. The former diagnosis faces the same issues as a standard descriptivist account, and the latter requires speakers to have shared views about what the role of the gender concept is, lest the disagreement be a merely verbal one turning on disagreements about the role of the gender concept. This, too, is not a plausible view.

Coordinating expressivism lets us make sense of these disputes in a way that requires neither a shared grasp of previously established semantic criteria for the relevant terms nor a shared grasp of the role played by the relevant concepts. Instead, we converge in the course of (successful) conversation on what attitude to associate with the utterances at issue. That means that people can disagree about how many genders there are, or what race someone is, without sharing background theories about what race and gender are or about the cognitive role of race and gender concepts. Since such disagreements patently do take place, coordinating expressivism offers a promising account of what’s going on in them.

Does Identity Politics Reinforce Oppression? An Answer from Social Metaphysics

Katherine Ritchie—Does Identity Politics Reinforce Oppression? An Answer from Social Metaphysics
City University of New York

Identity politics has been critiqued in a multitude of ways. One central problem, which I call the the Reinforcement Problem, claims that identity politics reinforces the existence of social groups rooted in bias, oppression, and disenfranchisement. Proponents of the problem claims that identity politics cannot be part of a successful liberatory project as it is paradoxical to seek to promote justice through reinforcing the existence of unjust social groups or identities. So, identity politics should be abandoned.

I argue that careful consideration of the metaphysics of social groups and of the practice of identity politics provide resources to dissolve the Reinforcement Problem. Identity politics involves the creation or transformation of groups in two ways. Neither succumb to the Reinforcement Problem.

First, identity politics can involve the formation of a new group based on intentional and overt commitments to values, projects, and political goals. The view fits with the way identity politics was originally argued for by the Combahee River Collective. It also coheres with arguments made in feminist philosophy. Using resources from social metaphysics offers a new way to conceptualize points that have been adhered to and proposed by scholars and activists.

Second, identity politics can involve the formation of new organized groups. These groups do not have identities based merely in oppressive power structures. Rather they are groups that are more akin to committees or teams. When combined with the transforming of feature groups, the creation of new organized groups also dissolves the Reinforcement Problem.

Social metaphysics has a role to play in vindicating identity politics from a central challenge. Identity politics can involve the transformation and creation of new feature and organized groups. It does not (at least when well practiced) involve the reinforcement or further reification of oppressive groups.

The Constitutive Aporia of the Relationship Between Individuals and Collectivity

Claudia Rosciglione—The Constitutive Aporia of the Relationship Between Individuals and Collectivity

University of Palermo

At which conditions are we allowed considering a plurality of individuals as a collective or social group? This is the question I address in reflecting on the relationship between the individuals and the collectivity considered in the process of the formation of the latter. In particular I take into account Margaret Gilbert's theory of the joint commitment and of the plural subject, on the one hand, and John Searle's theory of the collective intentionality, on the other one, to see the position they adopt regarding the tension internal to the expression "individuals as a group", in which the individuals are manifold and plural while the group is one and singular. Both theories aims at overcoming the tenet that the collective dimension is always reducible to the individual dimension and that the collectivity is a new subject overhead the individuals, which is endowed with a group mind. I show however that this common attempt at defining an anti-dualist third way between individualism and collectivism is submitted differently by these authors.

In Gilbert's view one is allowed speaking of a collective "We" if two or more subjects are jointly committed in a determined manner and a particular context. The joint commitment is realized when at least two individuals "are jointly committed to X as a body". I will dwell upon this image of the body because Margaret Gilbert draws the notion of the plural subject just from it. A group is actually as such because it is a plural subject and two or more individuals are one plural subject if they are committed to behaving as one single body. In this paper I argue however that in this view the irreducible dimension of the plurality is lost to an absolute unity that holds the individuals hostage, notwithstanding the theory of the plural subject. It's no coincidence that Gilbert refers to Hobbes' "real unity of all" to describe the social unity of the group intended as a plural subject.

Searle defines the "We" instead as collective intentionality that is a primitive concept in each individual. In this view one does not speak of creating a new subject such as the plural subject, rather of the cooperative, collective, intentional attitude that like all the intentional attitudes remains inside the single different individuals, which are distinct from one another. The group arises by bringing about all the collective intentionalities of the subjects who act by cooperating while preserving the plural dimension, because every one plays her part. This model, which Gilbert charges of subjectivism and individualism, has the merit of preserving the plurality of the individuals and of their mental states, which are not reduced to the absolute social unity that is identified with a new subject (plural). At the same time, Searle's view may seem to deflate the commitment to the strong social unity that should be constitutive of every group.

The aim of this paper is neither proposing a solution nor leaning towards one of these positions, rather arguing that these two views show the inherent aporia that is essential to the relationship between the individuals and the collectivity, which is constituted just in virtue of the aporia. The group is such if it is something open, hence continuously definable on the basis of the commitments and the intentions, which because of their having always an individual nature are compelling and at the same time reversible. Therefore the tension internal to the expression "individuals as a group"

cannot be eliminated because it is constitutive of the expression itself as well as of the relationship between individuals and collectivity that it denotes.

Modeling Group Agency and Social Norms Using Conditional Game Theory

Don Ross and Wynn Stirling—Modeling Group Agency and Social Norms Using Conditional Game Theory

University College Cork and Brigham Young University

If there are possible group agents then there must be possible group preferences. But the idea of a group preference raises a number of modeling challenges. There are major advantages in modeling the formation and maintenance of groups using game theory, not least the provision of a family of well understood solution concepts. However, if a group preference is simply an equilibrium of a non-cooperative game, then the concept of group preference does no work and is simply empty semantics. Team reasoning models agents as flipping, by gestalt-like processes, between alternative game models depending on whether players identify with their individual utility functions or with team utility functions. This device renders the group-forming mechanism exogenous to the game representations, so sidesteps the challenge instead of solving it.

Conditional game theory (CGT), uniquely among currently available modeling approaches, meets the challenge directly. CGT applies the syntax of Bayesian probability theory to games so as to model conditional preferences as formally analogous to conditional probabilities. Individual players' preferences are modeled as conditional in two senses, at separate stages of analysis: as subject to influence by other specific players, and as sensitive to the relative degrees of discord within groups that arises for different equilibria. CGT allows for formal identification of conditions under which, as propagation of social influence spreads through a group, and players modulate their preferences on the basis of other players' preferences, a group preference in a strictly defined sense may emerge. The theory allows us to derive a coordination ordering for a group which combines the conditional and unconditional (i.e., "classic") preferences of its members, in the same way as the joint probability of an event is determined by conditional and marginal probabilities. An operation of marginalization is defined that extracts players' ex post preferences once social influence has permeated the group, fully taking into account the social relationships and interdependencies identified in earlier stages of analysis. As these ex post preferences are unconditional, standard solution concepts such as dominance and Nash equilibrium (NE) can be applied to them, and we need not 'step outside of' endogenous strategic dynamics.

Marginalization can be recursively applied at the level of the group to determine whether a consensus choice exists; where it does, there is a group preference from which a group's best collective action profile can be derived. One gloss on the general technology is that it provides a model of socciation (dissociation), the extent to which, in aggregate, agents are (aren't) sensitive to one another's preferences. An agent's 'sensitivity' can be interpreted as meaning that she resolves some ex ante uncertainty about her own preferences through observing the preferences of others. We compare representations of norms that can be constructed using CGT, and show that Bicchieri's well-known analysis defines one member of a larger formal family of well definable concepts. We argue, finally, that it does not over-sell the power of CGT to say that it gives us formal conditions that must be satisfied when group agency is achieved.

Directed Duties, Shared Agency, and Wronging Another

Abe Roth—Directed Duties, Shared Agency, and Wronging Another
Ohio State University

Directed duties are essentially relational. A promise, for example, is made to a specific person, the promisee, and the obligation to keep it is directed or owed to the promisee. Other duties are monadic; they are fundamentally about the agent, and do not relate her to others, except perhaps incidentally. For example, perhaps I have a duty to refrain from destroying some pristine wilderness, or a duty to engage in occasional charity work. Though it might be in the interest of others that I fulfill such a duty, the duty itself doesn't relate me to specific others. I don't somehow owe it to anyone in particular to fulfill the duty.

Some explicate this directedness of obligation in terms of the distinction between acting wrongly and wronging someone – of there being someone in particular that you are wronging. To renege on a promise is, of course, to act wrongly. But it is, in addition, to wrong someone in particular – namely, the individual to whom one has made the promise. Relational wronging is, however, an imperfect gauge of the directedness of obligation. First, it turns out that violations of monadic duties can also engender relational wrongs. Assaulting a stranger is to wrong someone. But it does not involve a preexisting substantive relationship with the victim that we find in standard examples of directed obligation, such as those emerging from agreements or promises.

I argue that the directedness of promissory obligations is connected with the promisee's power to release the promisor from his obligation. She has an element of control over the obligation and thus over what the promisor does. This normative power and control over another is not restricted to explicitly moral cases such as that of promising. It's at this point that the issues raised by directed duties connect with shared agency. Some argue that when individuals act together, they are not only committed to a common end; they also have commitments to each other. This suggests understanding directed duties in terms of shared agency.

A shared agency model of directed duties is more promising than an account in terms of relational wronging. An advocate of the latter might, however, appeal to the law of torts to challenge the shared agency approach. Torts are harms one person visits on another for which one might be held legally liable. This can seem like the relational wronging. Moreover, theorizing about tort law and other parts of civil law was influential in the development of the notion of directed duties. So there is some plausibility to the thought that the law of torts will offer insight into directed duties. This poses a challenge to the shared agency model, because two complete strangers need not be related by contract, promissory interchange, or any form of joint enterprise for one to harm the other and to be subject to a duty (of compensation) directed to the other. I will show how a shared agency model of directed duty might respond to this challenge.

Institutional Identity

*Joshua Rust—Institutional Identity
Stetson University*

The English Crown is a circuitous institution that arguably finds its origins in Egbert's regime (802-839), the first king to rule over the whole of Anglo-Saxon England. What makes the reigns of Egbert and Elizabeth II stages of the same institution? Even a cursory understanding of English royal history might move one to simply reject this question: their reigns are not stages of the same institution because there does not seem to be a common enough thread, either in terms of constitutive rules, normative powers, or functions, that would strongly connect them. And yet we understand their reigns to be stages of the same institution, the Crown.

I argue that Derek Parfit's account of personal identity (1986) can be adapted to provide us with an account of institutional identity capable of accommodating the assumed continuity of the Crown.

Parfit famously sought to illuminate his account of personal identity by comparing a person to a club, nation, or political party. Just as a club can survive changes in membership, governing rules, and missions, Parfit suggested that radical changes in a person's psychological makeup are sometimes compatible with that person's survival. He also argued that under certain conditions, as when a club temporarily goes out of existence, the identity of these two clubs is indeterminate--there might not be a fact of the matter as to whether those clubs are the same club. Parfit appealed to institutions because he thought our intuitive understanding of a club's identity conditions are significantly less cloudy than the understanding of our own identity conditions: "even if we are not aware of this, we are naturally inclined to believe that our identity must always be determinate" (Parfit 1986, 217).

While our understanding of the identity conditions of clubs might be less cloudy it seems that some social ontologists are nevertheless tempted by some of the same distorting inclinations to which Parfit was responding. First, there is a tendency to think that there must be a common thread--whether articulated in terms of constitutive rules (Searle), normative powers (Hindriks), or functions (Guala)--that strongly connects the stages of every institution over time, just as Locke thought that a single psychological ingredient--shared memories--renders a person at t1 and a person at t2 the same person. I argue that such accounts cannot adequately explain why the reigns of Egbert and Elizabeth II are stages of the Crown. Second, some social ontologists preclude the possibility of genuine institutional indeterminacy. I argue that such indeterminacy represents a key context in which declarations of continuity play an important role in making an institution at t1 and an institution at t2 the same institution.

If Parfit could use our intuitions about clubs to help motivate his account of personal identity, which resists the idea that such identity requires a common psychological thread and is always determinate, then perhaps that account, in turn, might be reapplied to institutions, such as the Crown.

Interpersonally Scaffolded Affectivity: The Social Motivation Approach

*Mikko Salmela and Michiru Nagatsu—Interpersonally Scaffolded Affectivity: The Social Motivation Approach
University of Helsinki*

An influential approach in recent theorizing on collective and shared emotions analyses them within the Extended Mind framework as Extended Emotions (e.g. Krueger 2014, Krueger & Szanto 2016; Colombetti & Krueger 2015; Colombetti & Roberts 2015; León, Szanto & Zahavi 2017). The main idea is that, analogous to our use of tools for thinking, we construct affective niches with artefacts and other agents for regulating our feelings and emotions. A salient example is the handbag and its contents a woman carries (Colombetti & Krueger 2015). Slaby (2016) rightly criticizes the tendency of this approach to identify the individual as the autonomous ‘user’ of external environments as her affective ‘resource’, without questioning how individual affective and mental states are constituted and constrained by environmental scaffolding, both social and physical. We use Slaby’s critique as our starting point, but add two novel contributions. First, we argue that to make progress in this debate we need explicit analysis of the function and mechanisms of interpersonally scaffolded affectivity because the important affective scaffolding (including material one) often involves other agents or anticipation of their reactions (Griffiths & Scarantino 2009). Second, we provide such analysis, drawing on the social motivation hypothesis (Anonymized 2014). The social motivation is a particular psychological disposition whose main role is to orient humans toward affiliative stimuli, which yield social reward and enable the formation of social bonds. The function and mechanism of social motivation, we argue, are the key not only to understanding how interpersonally scaffolded affectivity works but to evaluating whether or not a particular affective niche promotes individual and collective well-being. In particular, we suggest that our orientation toward affiliative stimuli is so basic that it is not amenable to individuals’ self-oriented instrumental regulation but rather susceptible to systematic manipulation resulting in reduced sense of individual well-being. The social motivation hypothesis provides a framework for studying such problematic niches, e.g. Slaby’s example of emails as the “mind invasion” or the “corporate life hack”, as well as niches that are genuinely conducive to personal and collective well-being. Using this framework, we briefly review several empirical studies on face-to-face vs. online social interactions and discuss their general design implications for interpersonal affective niche construction.

Cooperation and Shared Intentional Actions

Jules Salomone—Cooperation and Shared Intentional Actions

City University of New York

In this paper, I attempt to locate cooperative activities among the phenomena captured by the literature focused on shared intentional activities. The seemingly paradoxical claim I make is that, pace Searle, Bratman and Ludwig, the study of shared intentional activities and that of cooperative activities should be kept relatively distinct.

At first, the concepts of cooperative activity and shared intentional action appear related. Suppose you and I share the intention to paint our house. If I act uncooperatively by knowingly buying your least favorite paint color when other similarly priced colors were available, I may undermine your sharing the intention to paint the house, especially if you intended for your least favorite color to not be used. This suggests shared intentional actions all involve cooperation in at least some minimal sense without which intentions could not be shared. But is cooperation in a more robust sense required for sharing intentions and thereby for shared intentional activities? Conversely, must parties to a projected activity share intentions for this activity to be cooperative? Combining these questions, are activities jointly intentional just in case they are cooperative?

First, not all cooperative activities require that their parties share relevant intentions. Cooperative activities require that parties share in some way their agential capacities, and shared intentions may well provide the adequate connection. But this connection condition can be met differently, e.g. through institutional roles and corresponding role-based obligations.

Second, not all shared intentional activities are cooperative. Two sets of reasons support this. On the one hand, I show that Bratman in his 1992 seminal article allows (albeit unintentionally) the possibility that noncooperative Nash Equilibria qualify as shared intentional actions. This finding is significant for the following reason. Among all recent accounts of shared agency, Bratman's would at first seem best suited to capture most (if not all) features of cooperative activities. In particular, his requirement that parties to a shared intentional activity intend other parties' relevant intentions be involved in the performance of that activity is apparently on the right track. Besides, Bratman explicitly intends his 1992 account to home in on shared intentional activities that are cooperative. The counterexample I offer not only shows that even under a robust conception of shared intentional activities, not all such activities are cooperative. It also suggests one should not expect a theorist to capture cooperation by adding even more stringent conditions on the sharing of intentions to those Bratman already defended. That is, one should not expect cooperative activities to be analyzed in terms of an exceptionally robust kind of shared intentional activity.

On the other hand, I argue we have good theoretical and ethical reasons to think that not all shared intentional activities are cooperative. For arguing so enables to discern distinctive injustices and pathologies from which parties to uncooperative, yet shared intentional, activities may suffer. In particular, I argue that the alienation and exploitation of some may well be partly explained by their sharing the intention to perform some activity that falls short of cooperation.

Adam Smith On Moral Sentiments in Structural Contexts

Kirun Sankaran—Adam Smith On Moral Sentiments in Structural Contexts
Brown University

Paradigmatically, we express blame and resentment towards other agents. In the Theory of Moral Sentiments, Adam Smith argues that resentment occurs as a natural reaction to a deficit in the quality of another's will towards us. Smith argues further that even when an action that harms someone need not reflect a bad quality of will, we might still be justified in blaming the causally-responsible agent. Smith gives a case in Book IV of the Wealth of Nations that seems (given the account in TMS) to license blame despite a lack of discernibly bad will. But curiously, he explicitly declines to do so. Instead, Smith argues that the rules under which the agents in question operate are the proper objects of blame and resentment. This creates two puzzles. First, there is a textual puzzle: Why does Smith decline to blame particular agents even if his theory explicitly provides the resources to do so? Second, there is a substantive puzzle. On Smith's theory, as on most others, proper targets of blame and resentment are agents. Non-agents are at best degenerate targets of those attitudes. And the thing Smith explicitly singles out--"the system of government"--is not an agent. It does not have anything resembling a quality of will to be assessed. So how do we make sense of blaming such a "social-structural" feature? In this paper, I answer the first question.

The textual puzzle is relatively easy to solve, but it makes solving the substantive puzzle more important for assessing the plausibility of Smith's overall theory. The explicit blame of the "system of government" and non-blame of particular agents serves a rhetorical purpose. The Wealth of Nations has explicitly reformist aims, and the passage allows Smith to emphasize the proper role of blame and resentment in this reformist project. I argue that while particular agents may be proper objects of blame--nowhere in the Wealth passage does Smith deny this--the excessive focus on blaming individuals distracts from the questions about social structure and its reforms that animate the book.

The substantive feature of Smith's account of blame that allows this reading to work is its reformist function. For Smith, blame and resentment are forward looking. Besides picking out a defect in the quality of another's will, they carry a demand that others modify their conduct in the future. I argue that this motivates Smith's argument in the Wealth passage. While there may be deficiencies in the qualities of particular agents' wills to be assessed, Smith's focus on institutional reform throughout the book suggest that he does not think particular agents are the right targets of the demand for better conduct in the future. It is clear that he thinks institutional structure--the "systems of government" that determine agents' incentives--is the right unit of analysis for social reform. And in keeping with his tendency to make moral claims in terms of sentiments, blaming institutional structures allows him to explicitly target them as objects of moral evaluation and reform.

Felicity Conditions of Document Acts

Gloria Sansò—Felicity Conditions of Document Acts

University of Turin

During recent years, philosophers such as Maurizio Ferraris and Barry Smith have maintained the importance of documents in acting within formal institutions (marriage, educational system, etc.). In these theories, some felicity conditions of document acts are mentioned, and some examples are provided. My aim is to make a contribution to this debate by presenting a more detailed and systematic analysis of these conditions. The starting point of my work are the felicity conditions of speech acts proposed by John Austin (1962). According to Austin, by means of utterances such as "I promise" or "I baptize", we perform an action. In order to avoid misfires, this action must be executed completely and correctly in accordance with a conventional procedure. In addition, in order to avoid abuses, this action must be executed with the proper intentions, thoughts and feelings. But unlike misfires, abuses do not have the power to make an act null. I argue that, with appropriate modifications and additions, some of Austin's insights can be suitable to account for document acts. In order to outline these changes, I focus on four peculiarities of document acts. First. Document acts are performed within formal institutions. Hence, it is more accurate to talk about laws, rather than conventions. The former provides that, in order to execute an act, it is necessary to produce a document and, sometimes, to utter some words. Second. Unlike the speech acts, document acts have a more evident material aspect. Paper documents have to fulfill some physical conditions, otherwise acts are null. Third. Austin claims that, in order to be successful, a speech act has to be grasped by all the participants. In document acts, this uptake is usually certified by actions such as signing or ticking boxes. However, from the fact that I have just signed a document does not follow that I have actually understood its content. Most importantly, some document acts do not even require this supposed uptake. If, for example, an order for payment is sent by certified mail, it is valid and binding even though the recipient forgets to pick up the letter at the post office, and hence he certainly ignores its content. Forth. According to Austin, in a felicitous speech act, intentions must be aligned with the act itself. If I promise to do X, I must have the intention to do X, and afterwards I must do X. A simulation entails an abuse. Austin, however, does not take into account the option of a bilateral simulation. In the case of a promise, if both parties simulate their consent, it is hard to say that the promise exists, and that simulation entails solely an abuse. In document acts, this issue seems to be less controversial. Suppose that, during the conclusion of a marriage contract, the involved parties simulate their consent. Without a declaration of invalidity, this act is still going to be valid and to exist. It raises a question. How important are intentions in document acts?

The Metaphysics and Politics of Personhood: Issues in the Social Ontology of Persons

Heidi Savage—The Metaphysics and Politics of Personhood: Issues in the Social Ontology of Persons
SUNY Geneseo

Many accounts of personal identity assume two things: (a) that selves should be integrated, and (b) that we can understand ourselves as psychological beings independent of our embedding environments. Our original assumptions (a) and (b) are intertwined in a particular way: the possibility of an integrated self depends upon whether a person can be understood independently of its embedding environment. I argue that both assumptions are mistaken for various reasons. The integrated self is a myth that can be sustained only for those who have adequate control over their environments, a luxury reserved only for the most privileged among us. For many of us, the integrated self is beyond our grasp, those of us who may have been raped, who live as outsiders within, who have mental health issues. All of these considerations raise the issue of whether an integrated self is a goal even worth pursuing. Understanding persons correctly requires rejecting (b) and replacing it with an alternative view on which we understand persons as fundamentally dependent upon their environments, a view that can accommodate fractured selves as full persons, and allows for various psychological configurations to satisfy the criteria for being a person.

Having a view that can accommodate those with fractured selves as persons is preferable to views that either require integration or that bind selves to their socio-political contexts. For example, one view of the self is composed of the idea that we construct our personal identities by constructing a coherent narrative about ourselves. While this view allows for more fluidity in what counts as psychological continuity than the standard theories of what constitutes psychological continuity, it nevertheless makes being a person depend upon having an integrated self, again, a condition not met by many who plausibly count as persons by anyone's measure. A replacement theory often put forward is the social constitution view. Persons should not be understood as having an intrinsic nature independent of their social environment. Selves are socially relational, and this relationality can accommodate various kinds of psychologies as satisfying the criteria for being a person. Still, the idea that selves are constituted by social relations may bind persons too closely to their particular socio-cultural context. However, a view like the one advocated for above, understands persons as dependent not only on social relations, but also upon other non-social relations to the environment. This has the advantage of allowing not only for various kinds of psychologies to satisfy the criteria for personhood, it also frees persons from particular socio-cultural contexts, a welcome alternative for those whose privilege does not allow them the luxury of an integrated self, and who yet need an avenue of resistance against their current socio-cultural context.

Communication, Inc.

Cosim Sayid—Communication, Inc.

City University of New York

Corporate bodies like joint-stock companies, legislatures, and philosophy departments communicate via publishing texts – sometimes contracts, press releases, or statutes – and how we ought to interpret what such corporate bodies meant via corporate acts of communication is an important problem. My focus is on the legislative case with ordinary statutes.

Two prominent approaches to dealing with this problem don't fare so well in the corporate context: actual intentionalism and purposivism. Although actual intentionalism – the idea that an agent's communicative intentions constitutively determine the content of what she communicates – works well in the individual case, corporate agents are comprised of many conflicting minds. So, even with an empirically adequate account in the individual case like the Gricean theory of communication, its explication in the corporate context would require either a group mind with corporate communicative intentions or a means of agglomerating individual communicative intentions into a corporate communicative intention. The former seems metaphysically hopeless, since it is hardly obvious that group minds are anything more than a metaphor. The latter is metaphysically suspect in the absence of an appropriate agglomerative principle.

Bratman's seminal work on shared agency might be thought promising as a potential agglomerative principle, but this is not actually so because there is no agglomeration from individual-level intentions, and it is moreover going to be false in many cases – such as the legislative one – that anything more was jointly agreed to besides some text being made law; it's not as though, for instance, legislators uniformly concur on the scope of quantifier-terms.

Purposivism founders in a similar way; legislators lack a common purpose in most cases, with the exception of enacting into law the text they voted on. Some legislators have an overriding purpose of pleasing their mistresses or party whips, others are keen to receive campaign donations for supporting some measure lobbyists favor, and so on. There are a motley of different purposes to consider even before one gets to challenges like Waldron's Wollheim machine.

Hypothetical intentionalism affords us objectivity of linguistic meaning by employing a non-actual intended interpreter. Although hypothetical intentionalism does not, and cannot, eliminate the grounds for interpretive disagreement in the way that actual intentionalism can in the case of individual agents, it nonetheless secures non-skeptical interpretive judgments in a manner consistent with the rule of law, while making clear the bounds of disagreement, which can be broached in other, sometimes more tentative ways, like creative interpretation, which is an exercise in circumscribed judicial lawmaking.

From the Social World to Stones and Back Again

Naomi Scheman—*From the Social World to Stones and Back Again*
University of Minnesota

Drawing on work on the social construction of kinds, I explore the social construction of particulars. An entity such as a university, musical performance, or (I have argued) an emotion, belief, or desire is socially constructed iff its being a complex particular depends on social practices in abstraction from which it disaggregates. Questions about what is or is not a part of the thing are answered by appeal not to physical relationships of contiguity, causality, etc., but rather to the relevant social practices.

This account of social construction relies on an important—albeit ineliminably vague—distinction between a complex particular and a jumble or heap. I explain the distinction by combining Quinean meta-ontology with radical explanatory pluralism: ontology follows explanation, rather than setting a priori constraints on it. We are ontologically committed to those things that figure in explanations we accept about how the world works, and we ought to—as we do in practice—accept many different sorts of explanations. Since our commitment to the explanatory value of our explanations admits of degrees, so does ontological commitment: some things are more robustly real than others.

The distinction between complex particular and jumble or heap is quite general, applying to stones vs piles of sand. What distinguishes a stone from a pile of sand is a matter of the degree to which each enters into explanations of how the world works, which in turn is a matter of the richness and robustness of the differences made to and consequently constitutive of the stone or the pile, and of the richness and robustness of the differences each has made to the things around it. The stone has a long narrative, interwoven with the narrative of other things in its vicinity: igneous, sedimentary, metamorphic; part of a larger formation, carried elsewhere by a glacier; capable of injuring toes that in turn move it as a single entity; capable of breaking windows. The pile of sand, by contrast, has little of what Spinoza called *conatus*: it came into existence randomly and has little propensity to continue hanging together, nor does it, as a pile, have much in the way of causal powers.

Knowing a real thing thus calls for attending to the narrative that constitutes its trajectory through time and space and checking in with its neighbors. On such a view of things and what it is to know them, humans and the social world are less discontinuous with the rest of the world than is commonly presumed. In particular, the non-human world is not indifferent: a stone is the thing it is because it responds to and affects some things and not others, and it responds to and affects those things in distinctive fashion, both because of the sorts of thing that it is and because of its particular history. We will better understand ourselves and the social world we create when we think of ourselves as distinctively enacting selective responsiveness—mattering, something that characterizes all the things that there are.

The Sources of 'We-Intention': Sellars, Durkheim, and Collingwood

Hans Bernhard Schmid—The Sources of 'We-Intention': Sellars, Durkheim, and Collingwood
University of Vienna

The proximate historical source for large parts of what is currently discussed under Searle's label 'collective intentionality' is Wilfrid Sellars' concept of 'we-intention'. But where did Sellars get the idea? Potential candidates abound, such as Jean-Jacques Rousseau's concept of the general will, or Émile Durkheim's concept of collective consciousness, or perhaps analysis of we-experience in early phenomenology.

Sellars' texts acknowledge no sources. We are thus forced to look at other clues such as biographical and conceptual proximity. Though it is possible that early phenomenological analyses influenced Sellars via his teacher Marvin Farber, and though Scheler's concept of co-experience can be construed as closely resembling Sellars' account of the 'we-mode', there is no evidence that there is indeed an influence. And Durkheim seemed to be off the table for lack of biographical connection and because in Sellars, there is no trace of the sulfuric smell of collectivism that has repelled so many of Durkheim's Anglo-Saxon readers. And Rousseau seems a bit remote, to say the least.

When the question of Sellars' source was examined in earnest about a decade ago, inquiries among Sellars' students seemed to provide another clou. Robin George Collingwood's concept of "social consciousness" was pointed out, which makes good biographical and conceptual sense. At the time Sellars' studied in Oxford, Collingwood was working on his *New Leviathan* (1947), where society is analyzed as "practical social consciousness" in terms of a "joint will". Though Sellars' we-intention is about values more than it is about society, Sellars' could have said about values what Collingwood says about society: "Society is the sharing of certain persons in a practical social consciousness verbally expressed in a formula like: 'We will go for this walk' or: 'We will sail this boat'. In this formula the word 'we' expresses the social character of the consciousness: the word 'will' its practical character. The concluding words are a definition of the common enterprise." (Collingwood 1947, 146). This seemed to close the case.

Upon recent discoveries, however, the case has to be reopened. Peter Olen and Stephen Turner (2015, 2017) brought Durkheim forcefully back into the picture, placing Durkheim at "the Origins of Collective Intentionality". To make their case, Olen and Turner present a smoking gun: A copy of Célestine Bouglé's *Evolution of Values* (1926) from the University of Pittsburg library. The book propagates Durkheimian ideas; it was translated into English by Wilfrid Sellars' mother, Helen Stalker Sellars. And it has an introduction written by Wilfrid Sellars' fater, Roy Wood Sellars. And: Wilfrid Sellars read it. The copy has Wilfrid Sellars' underlines in it.

It is thus now clear enough that contrary to what earlier 'historians' of collective intentionality analysis thought, worries concerning lack of biographical proximity should not rule out Durkheim as a potential source of Sellars' concept of we-intention. However, this piece of evidence in itself does, of course, not settle the conceptual question. Just how much of Durkheimian collective consciousness (or implicit discussion thereof) is there really in Sellars' we-intention? Are Olen and Turner right in claiming that Sellars' thinking on we-intention sprang out of Durkheimian roots? This talk examines Olen's and Turner's argument, places it in the context of other potential sources

of inspirations (which are not considered by Olen and Turner), and comes to a rather skeptical conclusion.

Inclusivity as a Virtue of Social Groups

Kate Schmidt—Inclusivity as a Virtue of Social Groups

Washington University in St. Louis

Epistemic wrongs are frequently understood as a sort of exclusion - agents are unfairly marginalized, dismissed as participants. One solution is to cultivate more inclusive epistemic norms. I argue that inclusion is morally valuable, and a property best ascribed to groups rather than individuals.

Epistemic injustice occurs when an agent is disrespected as a knower and an epistemic participant. Some philosophers have suggested that this type of wrong can be avoided by cultivating traits of justice, apt trusting, and open-mindedness. Inclusivity, the ability to include others in social activities, is a similar type of virtue that can be understood in contrast to epistemic injustice. People who are inclusive are able to fully listen to individuals who are otherwise marginalized. This can help to both prevent and repair epistemic wrongs.

Individual virtues are built around the motivations of the individual and are causally impactful due to the actions of the individual. I argue that some epistemic goods function at the level of the group. Inclusivity is a group level good that reflects a morally and epistemically valuable disposition to recognize and include marginalized individuals. Group level goods have at least some effects that are only seen at the group level. For inclusivity to count as a quality or achievement of a group, a sufficient portion of the individuals in the group need to possess appropriate dispositions and motivations. One way to see the dispositions within a group is by looking at social norms within a group. Groups that possess a valuable good can reliably implement the good - in this case inclusion.

Inclusivity is an important epistemic and moral virtue, it's best understood as a virtue that applies to groups rather than to individuals. A tendency to be inclusive cannot succeed based only on the dispositions of a single individual. Group norms have the potential to powerfully alter the inclusiveness of a group as a whole. The moral and epistemic benefits of inclusivity can only be understood and implemented at the group level.

Toward A Layered Account of Collective Intentionality

Michael Schmitz—Toward A Layered Account of Collective Intentionality
Universitaet Wien

Imagine a group of kids kicking a ball around, evolving certain patterns of play, and, over time, shared skills and a shared practice. Imagine further that they jointly form intentions to meet regularly, that they evolve a narrative, a body of beliefs about how their group and their game came about, and negotiate a set of rules for it, which is passed on in the oral tradition. Finally, imagine these rules are written down, various kinds of functionaries are appointed to enforce them and to organize various aspects of the sport, fixed procedures are put into place for electing these officials and for resolving disagreements within the organization, and so on. This vignette is meant to illustrate the idea that there are different levels or layers of collective intentionality. In my contribution I will explore continuities and discontinuities between these layers, offer a rough taxonomy and discuss relations between layers.

On all levels subjects act in certain modes in which they experience and represent each other as co-subjects of actional and perceptual, practical and theoretical positions towards the world. The levels differ in terms of the representational format of the relevant representational states, acts and artifacts. For purposes of orientation, I distinguish three broad levels / modes. The mode of joint action and attention has a nonconceptual representational format, the level of we-mode beliefs, intentions and group speech acts a conceptual and propositional one. The level of institutional reality can be characterized in terms of the essential role of writing and other forms of documentation and the fact that individuals and groups take positions in terms of their roles in institutional settings. For example, an individual might give an order or propose a plan for action in their role as head of FIFA. A more precise characterization of these layers and their relation is proposed in terms of such parameters as their degree of context dependence, the degree of the externalization and of the durability of the relevant representations and their degree of abstraction, the degree of differentiation of representational role, and the size of the relevant collectives.

Responsibility for Collective Inaction

David Schweikard—Responsibility for Collective Inaction

Europa-University Flensburg

How can individuals and groups be responsible for situations or outcomes they didn't actively bring about? In what sense can they be responsible for harm they could have prevented by means of a concerted effort? Questions of this sort point to scenarios in which no actions are performed, some which are nevertheless interesting to the theory of responsibility. It these latter cases of inaction that concern me in this paper. I explore the ways in which individuals and groups can bear specific sorts of responsibility for their inaction. I assume that in such cases, in which others suffer harm, there exists a duty of assistance on the part of all those who could have actively prevented the harm or who could have taken steps to prevent it. And I take it to be plausible to say that if preventing a harm requires some form of collective effort, implying that no individual could perform the respective action on their own, then the duty of assistance is also in some specifiable sense collective.

Having thus set the stage, I then first address cases of inaction that more directly involve kinds of collective omission. Here I refer to shallow-pond-type cases in which, say, one individual ("A") is in danger of drowning and can only be saved by the effort of three agents who are nearby and jointly capable of performing the required rescue manoeuvre; and in which the manoeuvre is not performed. What kind of responsibility we can correctly ascribe to whom depends on what went wrong. This again we can arguably only say by reference to an account of what it takes to perform a joint action such as the one required. The account I favour – here sketched in idealized fashion – requires that, for all three agents in question, they are each and jointly aware (1) of A's situation, (2) of their joint ability to act, (3) of their duty to act, that they (4) coordinate to decide to act and (5) that they act. Their responsibility for inaction can then be characterized as due to a failure to form specific joint attitudes (awareness and intention) and a failure to perform the required action.

In a second step, I turn to cases in which the harm an individual ("B") suffers cannot directly be attributed to individual or collective action or inaction, but is instead due to structural features of B's situation. Here I have in mind cases like Sandy's (described in I.M. Young's *Responsibility for Justice*, chap. 2), in which it is impossible to trace the injustice she suffers linearly to some individual or collective agent. I want to argue (in Young's spirit) that by not actively changing social structures that are harmful to some, by being inactive, all those who contribute to and uphold them share responsibility for the harm these structures cause. I go on to sketch how the distinctions introduced above help to elucidate the forms of responsibility this implies.

Collective Inaction and Group-Based Ignorance

Anne Schwenkenbecher—Collective Inaction and Group-Based Ignorance

Murdoch University

This paper discusses morally wrongful collective inaction and the problem of group-based ignorance. Some of the many things that we could do together with others but fail to do are morally wrongful inactions. While the list of our – individual and collective – non-actions is infinite, not everything that I (or we) fail to do is some form of inaction that is plausibly attributable to me (or us). ‘Collective inaction’ is the unintended failure of two or more agents to perform a collective action or produce a joint outcome where that action or outcome was collectively feasible and where the individual agents had group-based reasons to perform (or produce) it. In a second step, I will discuss the role that ignorance plays in excusing morally wrongful collective inaction. I identify three different kinds of collective knowledge (common, pooled, or public) and corresponding types of group-based ignorance. I conclude by showing that inaction is excusable where ignorance sufficiently weakens agents’ group-based reasons for action.

Do Group Agents Resemble Psychopaths? (And If So What Does That Mean for Collective Moral Responsibility?)

*Anne Schwenkenbecher—Do Group Agents Resemble Psychopaths?
Murdoch University*

Many philosophers think that groups can be considered moral agents of sorts, provided they are structured in the right way. Corporations and states are amongst the most common examples for moral group agents. While scholars argue that these kinds of groups can be rational agents, there are very few who think that we can speak in a meaningful way about group agents having emotions. In that regard group agents seem to resemble human psychopaths, who are usually thought to lack important emotional capacities. Where human agents are concerned, the lack of empathy in particular is regularly considered to diminish an agent's capacity for moral reasoning and their moral responsibility. Nonetheless, we tend to ascribe moral responsibility to group agents such as corporations. Should this discrepancy worry us? Should it encourage us to be more open to arguments about group agents' emotions? Or else, does it undermine the idea of group responsibility? I will suggest that for a group to enjoy the epistemic advantages of moral emotions it is sufficient that its members have the right kind of shared, interlocking emotions. As such, group agents can be fully-fledged moral agents in the absence of group-level emotions.

How Is This Stuff Fundamental?

*Keyvan Shafiei—How Is This Stuff Fundamental?
Georgetown University Philosophy Department*

What is metaphysically fundamental? Some philosophers claim that metaphysical inquiry should study the categories that carve reality "perfectly at the joints" (Sider 2011, 6). Others claim that metaphysics should move beyond the fundamental. Elizabeth Barnes (2014), for example, argues that debates about metaphysical fundamentality reflect a kind of philosophical parochialism that hinder research into issues that can help us make better sense of our everyday practices. This includes substantive debates about what race and gender are. If she is right, then the insistence that metaphysics should focus on questions about fundamentality will limit the use of metaphysical resources to a narrow range of questions, leaving many significant features of our world unexplained. I think Barnes's worries are legitimate. I worry, however, that this dispute is fundamentally (sorry!) stuck in the mud. This way of organizing the dispute makes it hard to seek common ground. Thus, the study of metaphysics has fissured in a way that might foreclose possibilities for productive future collaboration between self-avowed fundamentalists and non-fundamentalists.

In this paper, I argue that the dispute over fundamentality should be reoriented toward a discussion of what we mean by 'fundamental'. First, I build on Barnes's suggestion (2014) that what is fundamental may extend beyond the natural. I define 'fundamental' by reference to the 'stuff that centrally give meaning to our interactions with the world'. In light of this argument, I argue that our ways of interacting with fundamental features of the world are mediated by the social structures that the non-fundamentalists want to metaphysically make sense of. In other words, I argue that both the structure of the world and the social structures that mediate our experiences of the world are fundamental. Building on Kristie Dotson's (2012) recent arguments, I will then caution philosophers against boundary policing in the context of disputes over what is and is not metaphysically fundamental. I argue that our understanding of what is fundamental should be dynamic, and we should be deeply receptive to the possibility that what we take to be fundamental extends beyond the natural.

Group Action Without Group Minds

Kenneth Silver—Group Action Without Group Minds

University of Southern California

We frequently speak as if there are groups and as if those groups can do things. One set of questions concerns whether groups really exist and how best to understand their existence (as well as the differences between groups, collectives, corporations, etc.). Another kind of question, however, concerns whether the behavior of groups can appropriately be construed as actions. Here, I focus entirely on the latter question, and I give a novel view of group action.

To show that group behavior amounts to action, it would be best to fit their behavior into a general account of action. However, nearly every account of action from the philosophy of action requires that the agent have mental states such as beliefs, desires, and intentions. And unfortunately, theorists are split over whether groups can instantiate these states - typically depending on whether or they are willing to accept functionalism about the mind. Luckily, we can avoid this debate.

Crucial to most view of action is that the agent's actions are appropriately explained by or done in light of her reasons. The agent is sensitive to her reasons; she is answerable for her conduct and answers in terms of her reasons. Given this, I provide a general view of action in terms of terms of responsiveness to one's reasons. It may be that creatures like us are not able to satisfy this account without mental states to mediate our reasons, but I think that groups can. The group's members (or their mental states) can fulfill whatever role mental states typically play in our actions. The reasons central to my account are motivating reasons, so I consider several accounts of motivating reasons and argue that groups can have them (and would even without mental states).

If the arguments are successful up to this point, then we will have a way of characterizing the behavior of many groups as actions. Though this is the primary objective of the paper, many of the people thinking about this topic are more interested in the issue of group or collective responsibility. The question there is whether it is appropriate to hold groups morally responsible for what they do, or whether all of the responsibility of a group redounds to the members. In the final section, I consider how the view of group action I advance bears on this question. One worry is that by accepting a view that does not require groups to have mental states, we have forfeited something necessary for them to be morally responsible. I show that this will not be the case on several different views of moral responsibility. This does not demonstrate that groups are morally responsible, and I acknowledge what remains to be shown in order to establish group responsibility. Still, I take defending a view of group action to be an important step.

The Cambridge View of Emergence

Yannick Slade-Caffarel—The Cambridge View of Emergence

Cambridge Social Ontology Group, King's College London, Sciences Po Paris

This paper will outline and defend the theory of emergence developed by the Cambridge Social Ontology Group and largely published within the works of Tony Lawson. In contrast to authors such as John Searle who argue that any notion of emergence necessarily implies causal reduction, the Cambridge view argues that in all cases examined, and indeed in all cases examined within the social realm, that emergents are at once causally and ontologically irreducible and, therefore, arguments supporting theories of causal reduction or downward causation (also known as top-down causation) are equally proscribed.

The Cambridge theory argues that the emergent nature of a novel totality's relational organising structure is largely ignored within the literature and that this has led to erroneous claims regarding the prevalence of causal reduction and downward causation. In Cambridge Social Ontology, the novel totality and its organising structure emerge simultaneously with both having causal powers and an existence that are not reducible to their disorganised constituent components. Moreover, once it is understood that the novel whole, its organising structure, and the components so organised all exist at a higher level of reality, then it no longer makes sense to speak of downward causation. A whole acts through its parts, not upon them.

The Cambridge theory of emergence is crucial to Cambridge social ontology as it at once grounds the conception as thoroughly naturalistic—indeed the Cambridge view argues that the emergence of novel totalities constituted by components at a lower level of reality is characteristic of all reality, from quantum fields to complex social phenomena—while also providing justification for the study of social science, and an argument for why such study can indeed be scientific in the same sense as the natural sciences.

This paper will argue that these implications of the Cambridge theory of emergence are necessary elements of any conception of social ontology. For a theory to hold, I will argue that it must at once be able to show that it is consistent with our best knowledge of how the world exists, that the study of society is justified due to an inability to understand it through simply studying some lower-level phenomena, and that such study can be scientific.

Reliance as a Source of Shared Action

Matthew Smith—Reliance as a Source of Shared Action
Northeastern University

Many theories of the social presume that even the weakest forms of collective action require agents intending to act jointly or at least believing that they are acting jointly. On these views, there must either be some kind of shared understanding of the joint activity or some authoritative source that at least most participants are disposed to agree determines the nature of the joint activity. The systematicity that unifies disparate actions into collective or joint action is highly intellectualized.

But, theories of structural injustice (as well as accounts of distributed mass political action) seem to require some form of collective action in the absence of the robust forms of intellectualized systematicity at the core of many theories of shared action. In cases of structural injustice, in particular, participants whose actions constitute the objectionable institution often flatly deny that there even is a such a structure. For example, the act of claiming White privilege reproduces the institutions that make that privilege possible. But, White people often are not even aware that they are claiming that privilege. Often, those who most strongly deny that there are racist institutions are the ones whose actions play crucial roles in constituting and reproducing those institutions.

This paper aims to make sense of this phenomena without holding that structural injustice is a mere unintended consequence of individual actions, and without the elimination of collective action in favor of something less demanding, such as collective responsibility.

The paper attempts this by showing how the human capacity for action - agency - is realized not merely through intellectual abilities but also in non-mental structures. This is because all actions involve reliance on the world, and all successful actions involve successful reliance on the world. In order to raise one's hand, one must be able to rely on one's body properly functioning. In particular, if one's arm wasn't working - suppose it had been paralyzed by an anesthetic - then the action of raising one's hand is not be possible. Failures of reliance are failed actions.

Now consider more complicated actions, such as cooking dinner. To cook dinner, one relies not only on one's body properly functioning, but also on the stove properly functioning. If one's stove malfunctioned, then one could not cook. This extends beyond material technologies. For, many systems on which we rely when we act are socially constituted. Action that necessarily involves reliance on these systems are irreducibly social. For example, language necessarily involves reliance on at least semantic and syntactic conventions. Thus, speech acts are necessarily social. That is, all communicative acts are instances of some form of collective agency.

The paper argues that this point applies equally to actions that necessarily involve socially constituted affordances such as White or male privilege, along with other actions that necessarily involve reliance on unjust institutions. We conclude by arguing that an agent's lack of awareness of these structures does not render actions realized through reliance on institutions any less social. Rather, it simply reveals that the agent is ignorant (on at least epistemic grounds) of the nature of her action. It follows that I must deny a strong form of practical knowledge, but I show that I can still accept a weak but attractive form of practical knowledge.

The Reframing of Social Choice Theory Under the Consideration of Conditional Self-Interest

Sabine Spangenberg—The Reframing of Social Choice Theory Under the Consideration of Conditional Self-Interest

Richmond University

Social ontology aims to enlighten our understanding of social phenomena. The observation and explanation of social phenomena is a core objective of social science of which economics serves as a sub-discipline. The modelling of behavioural relationships mostly employs the approach of factual empirical observation and conversion into a theoretical model to be applied to and tested in a suitable scenario. Much focus is placed on value free judgements in economics, aiming to make this social science a positive one. To achieve this, economists commonly use assumptions which further allows the observed facts to be converted into mathematical functions. Modern social choice theory is a core subject area of public economics with the outlook to engage in welfare considerations and the understanding of the functioning of policy tools. Whereas the functioning of many economic instruments considers the individual economic agent as the central behavioural focus of the investigation, assumptions are made about the identification of social preferences. Modern social choice theory places emphasis on democratic traditions and infers preferred social states from individual preferences implying also that individuals' utilities can be aggregated into collective utility. This was refuted by Kenneth Arrow's impossibility theorem imploring that no method existed that allowed an aggregation of individual preferences into an ordered set of social orders. This paper considers the implications of alternative definitions of social values and there with the understanding of welfare. The so far ignored conditional self-interest as suggested by Cicero will be juxtaposed to Hutcheson's consideration of human moral sense and affection. This will inform further on the limitations of utilitarian assumptions and highlight some fallacies of the otherwise brilliant Arrow-Sen formulation. A kinder approach to human nature that also considers the existence of thought for others generates varied assumptions and an alternative understanding of social phenomena. A de-individualised redefinition of social states is necessary to arrive at a true set of options.

Knowledge, Behavior, and Rationalizability

*Todd Stambaugh and Rohit Parikh—Knowledge, Behavior, and Rationalizability
City University of New York and CUNY Graduate Center, Brooklyn College*

In social situations, rational agents base their actions on knowledge and beliefs about the circumstances at hand. At the very least, this includes knowledge about the possible strategies of themselves and other agents and the preferences agents have regarding the outcomes of those strategies, but often those outcomes also rely in part on some other facts of the world beyond the control of those involved. For example, if the weather is nice, two people might prefer to meet at a park, but otherwise a library would be preferable. With this in mind, it would be irrational for one of them to go to the park knowing that it will rain. Without that knowledge, even if it will in fact rain, going to the park might still be considered a rational action. Put simply, the concept of rationality depends in part on a coherent notion of knowledge.

Bernheim (1984) and Pearce (1984) each defined the broad game theoretic solution concept of rationalizability (a generalization of Nash equilibrium), which is built on the premise that rational agents will only take actions that are a best response to some strategies that they consider possible for the other agents to use. This accounts for the consideration that the other agents are rational as well, which limits the possible actions to which a particular agent must respond, enabling further elimination of strategies until the set stabilizes. This process is fundamentally built on the strategic considerations that underlie behavior, but does not address the epistemic considerations described above.

It is natural then to generalize the notion of rationalizability for games in which payoffs depend not only on the strategies of the players, but on some facts of the world as well. This will enable examination of the interplay between the traditional strategic considerations and those which are the product of some state of knowledge. An account is presented of what it means for an action to be rationalizable relative to a particular state of affairs, and in turn relative to a state of knowledge.

In order to create a situation in which the state of knowledge is better suited to their preferences, social agents communicate with one another. Questions regarding the states of knowledge obtained through communication have been explored by Halpern & Moses (1990), and Sietsma & van Eijck (2011) among others.

Additionally, a class of games is introduced called Epistemic Messaging Games (EMG) which involve a period of communication among the players prior to the play of the game. This communication stage results in a particular epistemic state which can be quite complex, and thus requires a broad framework to examine. To this end we use a version of the kind of history based models presented by Chandy & Misra (1986) and Parikh & Ramanujam (2003), which frames individual knowledge as a local projection of a global history. Using this technique, we present a general account of rationalizability for certain subclasses of EMG designed to model various kinds of social situations.

Gaslighting: Individual and Collective, Epistemic and Manipulative

Cynthia Stark—Gaslighting: Individual and Collective, Epistemic and Manipulative
University of Utah

My paper develops an account of what I call “collective manipulative gaslighting (CMG)”. This is a form of manipulation, perpetrated at the level of the collective imagination, that induces members of certain groups to see their accurate perceptions as groundless and to therefore not take themselves seriously as interlocutors. As such, it is a type of psychological oppression. My essay has four parts. The first distinguishes CMG from epistemic gaslighting and from individual level gaslighting. Epistemic gaslighting is a type of testimonial injustice where a person’s claim to have been harmed is not believed due to a credibility deficit she suffers on account of her social identity. Manipulative gaslighting may include a failure to believe a person’s testimony, but more generally is accomplished by the perpetrator of a harm positioning himself as the victim when confronted about his wrong-doing. Epistemic and manipulative gaslighting have been primarily theorized at the level of individual interaction. However, both occur en masse, as when victimized women are suspected of making false accusations of rape against celebrities or when they are seen as ruining their celebrity rapists’ lives by bringing charges against them. The second part of my essay explains how the tactics used by manipulative gaslighters in individual interactions can be elevated to the collective level. These include “targeting”—exploiting a person’s feelings of inadequacy and “diverting”—accusing the victim of exaggeration or distortion and then using her legitimate anger at this as evidence of that she is the perpetrator. The third section of the paper explains why CMG is not merely a case of reasonable disagreement or defensiveness. The fourth examines two examples of CMG. One is from the philosophy blogosphere and the other is captured by Manne’s account of “himpathy” in her recent book *The Logic of Misogyny*.

Joint Action Without Theory of Mind

Daniel Story—Joint Action Without Theory of Mind

University of California, Santa Barbara

This paper is motivated by a tension between, on the one hand, leading philosophical accounts of joint action which presuppose that participants have theory of mind and, on the other hand, the idea from developmental psychology that joint action is importantly developmentally prior to theory of mind. First, I argue that, in light of this tension, we need an account of joint action that does not require participants to possess theory of mind. Second, I examine and reject two attempts to give such an account. Stephen Butterfill (2012) proposes a novel form of joint action that does not require shared intentions, but I argue that his view is extensionally inadequate and unparsimonious. Elisabeth Pacherie (2013) contends that individuals can share an intention via a special form of reasoning without explicitly representing others' intentions or beliefs, but I argue that this special form of reasoning itself presupposes theory of mind.

After rejecting previous attempts, I proffer my own account of joint action that does not require participants have a theory of mind. I draw upon a thesis due to Abraham Roth called Practical Intimacy (Roth 2004), according to which it is possible for an agent to take up and act on another agent's intention without independent deliberation. On this view, intentions can be transmitted or "preserved" interpersonally in much the same way that intentional states are preserved intrapersonally via memory (cf. Burge 1993). I argue that Practical Intimacy gives us a straightforward way of understanding how agents without theory of mind can share intentions and act jointly. Basically, I contend that agents without a theory of mind can possess intentions with a joint content that are transmitted to them by agents who are more socially-competent via the mechanisms that underpin Practical Intimacy because these mechanisms do not involve representing mental states as such. If this joint content is a joint action involving the issuer and the addressee and both act on that intention, then it is literally true that the two act on a shared intention. In many cases, this will be sufficient for joint action. This model of joint action with young children comports with the idea, widespread in developmental psychology, that children are only able to participate in social activities that are structured by more competent partners (Brownell 2011). Whereas previous accounts of joint action without theory of mind put the burden of participation entirely on the child, the version I proffer more-or-less outsources the cognitive demands of joint action through the mechanisms of Practical Intimacy. For this reason, I conclude that my view not only gives an account of joint action without theory of mind, but it also comports better with the relevant empirical evidence.

Toward a Behavioral Account of Social Institutions

Megan Stotts—Toward A Behavioral Account of Social Institutions
McMaster University

It is a commonplace observation that social institutions are deeply dependent on humans. For instance, the government of Brazil, the Catholic Church, and the Japanese language depend on humans for their existence in a way in which trees, oceans, and solar systems do not. A well-established view in the literature is that, in particular, the existence of social institutions depends on collective acceptance—that is, on our shared mental states. In contrast, I will argue that we should eschew mental states of all kinds and provide an entirely behavioral account of social institutions. This view respects the insight that social institutions are deeply dependent on humans, but with a twist: social institutions depend just on what we do, not on what we think.

I'll begin by arguing against the collective acceptance approach to social institutions, by gathering together a series of counterexamples already in the literature that illustrate the possibility of social institutions in the absence of collective acceptance. A natural response to the problems that the collective acceptance approach faces is to turn from collective to individual mental states. However, I contend that this sort of approach won't work, either, because there are still counterexamples: cases in which a social institution exists whose nature seems to contravene the individual mental states of its participants.

In my view, the way forward is illuminated when we turn our attention to non-institutional, merely social phenomena among other species, such as ant colonies. Here we can see that it is possible for social phenomena, at least, to exist without any mental states at all. My suggestion is that we should take our cue from these merely social phenomena and seek an entirely behavioral account of social institutions. Specifically, my proposal is to ground social institutions in copied behavior that clusters into roles, where there are equally accessible alternatives to some of that copied behavior, and where the behavior works together to promote some result(s). Importantly, this approach still allows for a distinction between full-blown institutional phenomena and merely social phenomena (such as ant colonies) by noting that the behavior that gives rise to merely social phenomena lacks the features of copying and alternatives.

An important consequence of this behavioral approach to social institutions is its implication that despite what we may assume or hope, the nature of our social institutions is not something we can directly determine by thinking of them in a certain way. For instance, a social institution that many (or even all) of its participants think of as promoting justice may actually turn out to be an institution whose purpose is oppression of a particular group of people, as a result of features of the behavior in which participants actually engage. Additionally, by taking the focus away from how participants think of their social institutions, we also gain the ability to easily explain how individuals can participate in a single institution together while having radically divergent beliefs, desires, and intentions about the nature and purpose of that institution.

Asymmetric Cases of Joint Actions

Anna Strasser—Asymmetric Cases of Joint Actions
Humboldt-Universität zu Berlin

Joint actions are a paradigmatic case of collective intentionality and group cognition, but many characterizations implicitly assume that all participants have more or less the same socio-cognitive abilities. Moreover, standard philosophical notions of joint actions seem to be tailored to human adults only (Bratman 2014).

Focusing on mixed groups with an unequal distribution of socio-cognitive abilities, I argue for two claims: First, not only human adults but also infants, non-human animals, and artificial agents are able to act jointly. Second, joint actions are possible with an asymmetric distribution of socio-cognitive abilities and this requires further explorations of topics such as responsibilities.

According to developmental psychology, children are socially interacting beings and engage in joint actions even though they do not fulfill the demanding conditions required by a standard philosophical conception. In this paper, I focus on a controversial case, namely joint actions of mixed groups consisting of human beings and artificial systems.

If artificial agents qualify as proper social agents in joint actions we are confronted with questions about what duties and rights they deserve as interaction partners. To lay the foundation for such discussions I develop an extended conceptual framework of joint actions clarifying minimal necessary conditions while leaving space for asymmetric cases.

In order to overcome the tendency to restrict socio-cognitive notions to sophisticated living beings, I refer to recent minimal approaches (Butterfill & Apperly 2013; Michael et al. 2016), which present a promising starting point for developing a broader framework. By questioning the necessity of conditions which come with standard conceptions they suggest so-called minimal versions of socio-cognitive notions. For example, the notion of 'minimal mindreading' (Butterfill & Apperly 2013) specifies the very minimal presuppositions of how agents can anticipate the behavior of others by questioning the necessity of overly demanding cognitive resources, such as the ability to represent a full range of complex mental states and a mastery of language. By requiring less demanding conditions this notion can explain success in mindreading tasks of various types of agents. Most significantly, with regard to artificial agents, it does not require conscious reasoning.

For the aim of this paper, I encourage the rationale of questioning seemingly necessary conditions in order to develop a notion of joint action which can account for artificial agents. I show that some conditions introduced by the standard notion rather serve as biological constraints specific to human beings and should not be used to exclude artificial agents from the start. Thereby, the space of group cognition can be expanded from living beings to non-living agents. After specifying under which circumstances artificial agents are able to act and coordinate, namely to understand other agents, to reciprocally exchange social cues and rely on the commitment of other agents, I explore the consequences such asymmetric cases have for our understanding of responsibility. The aim of this paper is to integrate minimal and standard notions in one framework so that we can do justice to asymmetric cases in future debates.

Group Agency and Homuncularism

*David Strohmaier—Group Agency and Homuncularism
The University of Sheffield*

Theories of group agency defend the attribution of minds and propositional attitudes to groups (e.g. List & Pettit 2011, Rovane 2014, Huebner 2014, Tollefsen 2015, Epstein 2015, 2017). A theory of group agency might, for example, ascribe to a philosophy department the intention to hire a new professor. My paper addresses two tightly related questions: Can we apply a version of the homuncularism objection to theories of group agency? And how strong is this version of the homuncularism objection?

Theories of group agency are typically homuncular, that is they imply an ontology which postulates mentality at multiple nested levels, since human groups with minds have members with minds of their own. There is little point in denying that candidates for group agency like a philosophy department contain minded agents, that is homunculi.

Postulating cognition at multiple nested levels, however, is generally seen as a problematic move. My paper surveys the associated worries and argues that only one version of the homuncularism objection applies to theories of group agency. This version is not the most common formulation and has only pro-tanto strength. It can be formulated in one sentence: To postulate mentality at the higher group-level, while already having mentality on the lower individual-level, adversely affects the ratio of explanatory benefits and costs.

Having identified the homuncularism objection as it pertains to theories of group agency, I assess its strength and, in particular, whether it can justify any universal principles restricting the attribution of group agency. For this purpose, I draw on work by Bryce Huebner (2014), who has postulated just such three universal principles. These principles are motivated by the homuncularism objection and are supposed to guarantee that we only postulate group agents which bring sufficient explanatory benefits.

I argue that group agents violating Huebner's principles can nonetheless contribute sufficient explanatory advantage. For this argument, I use an example from rational choice sociology (Goldstone 1994, see Risjord 2014): Revolutions are hard to explain within a rational choice framework since a massive free-rider problem arises. For an individual participating in a revolution is dangerous while their contribution is highly unlikely to make a difference. Introducing group agents solves the problem. As Goldstone writes, "if we focus on groups as the key actors that precipitate revolution, [...] we may find many circumstances in which it is rational for groups to undertake revolutionary action" (Goldstone 1994: 140). In this instance, group agency brings explanatory benefits in violation of Huebner's principles.

While I use Huebner's principles for illustration, the results concern the general strength of the homuncularism worry as pertaining to group agency. In virtue of its pro-tanto nature, the homuncularism objection cannot motivate principles as strong as those offered by Huebner. Unexpected explanatory advantages of group agents, such as in the case of revolutions, can outweigh the homuncularism worry. I conclude that the homuncularism worry has little general force against group agency attributions and can only undermine such attributions in specific cases and with regard to empirical research.

Habitus Contra Background: Clarifying the Role of Socialization and Cognition for Social Ontology

*Richard Strong—Clarifying the Role of Socialization and Cognition for Social Ontology
Villanova University*

The aim of this paper is to promote the concepts of habitus, capital, and fields/social space as developed by Pierre Bourdieu as a promising framework for social ontology (Searle 1995; Haslanger 2012). Bourdieu's ensemble of interrelated concepts has the potential to clarify a number of thorny issues in social ontology such as vertical relations of power, asymmetrical determinations of value and meaning, social reproduction and change, various scales and types of social aggregation, human kinds and classifications, and the welter of problems associated with the cognitive and practical grounding of social reality in the minds and bodies of agents in transactions with other people and things. Here, I focus primarily on the virtues of the concept of habitus over the Searlian notion of the background (Carmen 2001; Mallon 2016).

Habitus refers to the sociogenetically acquired impersonal schemas and systems of dispositions that function as the principle of practical action, perception, and apperception (Bourdieu; 1977). Bourdieu's conceptual apparatus, developed, refined, and tested through fieldwork, is designed to capture the social world and the basic constituents of social reality without falling prey to a number of distorting oppositions such as qualitative phenomenology opposed to quantitative social physics, social structures (e.g., collective intentionality) opposed to individual agencies, explanations from causes opposed to explanations from reasons, and agency opposed to structure. Habitus is a concept that is intended to escape such dilemmas and many more while also providing a positive ballast for theorizing social reality (Wacquant 2014).

John Searle explicitly acknowledges an affinity between his articulation of the background and Bourdieu's concept of habitus (Searle 1995). Such similarities are largely superficial but provide a clarificatory point of contrast between these two divergent orientations to the social world. First, whereas Searle grounds social reality in uniform collective intentions Bourdieu stresses the differential and relational constitution of social reality which can further serve as the basis for understanding relative degrees of power (vertical social relations) as well as inter and intra-class and intra and inter-group variations and concordances with respect to social meanings and values. Second, whereas Searle speaks of societies and social reality at presumably very large scales Bourdieu's concept of habitus enables us to better grasp social ontology at different scales (e.g., village, institutions, and nations) and across different types of social aggregation (e.g., classes, teams, large organizations, and specialist microcosms). Third, whereas Searle posits singular status-functions Bourdieu foregrounds the polysemous and often fuzzy nature of socially recognized meanings and values which are cognitive, conative, and affective (Searle 1995; Bourdieu 1979; Taylor 1997). Fourth, whereas Searle writes of social reality as constrained by freedom and rationality Bourdieu speaks in terms of the arbitrariness of much (not all) of the social world; the entrenched and codified givens of the social world are largely arbitrary – stemming neither from reason nor a natural social order (Bourdieu 1990). Fifth and finally, whereas Searle speaks of a presently inscrutable neurophysiological background, Bourdieu puts forth a falsifiable social psychological model of habitus (Marcoulatos 2003; Lizardo 2009; Theiner & Fogle forthcoming).

Social Categories, Social Construction, and the Problem of Passing

Ásta Sveinsdottir—Social Construction, and the Problem of Passing
San Francisco State University

Social constructionists about race, gender, and other contested categories face a prima facie problem with the phenomenon of passing: if to be a woman, for example, is simply to occupy the position of being a woman, how can we make sense of the claim that some people are passing as women when they aren't really women. Or, to use another example, how can we make sense of the claim that some people pass as being of a race that they aren't really? Does the social constructionist have to say that all talk of really being F is just in error? What other options are there?

I will address how a conferralist can address the problem of passing by drawing on the work of Ásta in *Categories We Live By* (Oxford UP 2018).

I argue that the conferralist cannot claim that someone passes as a woman when they really aren't, but that the conferralist still has a way of making sense of the intuition that there is some misgendering or another sort of misjudgment taking place. On the conferralist picture the subjects doing the conferring are attempting to track a base property and their judgment as to whether the person in question has the base property is what confers the social status of being a woman. I maintain that the conferralist has to say that although a person does not pass as being a woman, they pass as having the base property for the conferral for being a woman in the particular context. This is enough to make sense of the phenomenon of passing in the sense that with that distinction we can make sense of the affective responses people have at the discovery that the person wasn't quite what they thought they were.

Responsibility as a Relational Property

Larisa Svirsky—*Responsibility as a Relational Property*
The University of North Carolina at Chapel Hill

Philosophical views of responsibility often identify responsible agency with capacities like rationality and self-control. Yet in ordinary life, we hold agents who have significant impairments in these capacities (i.e., “marginal agents”) responsible all the time. Consider the parent who holds her five-year-old responsible for not screaming in a restaurant simply because he is bored, or the therapist who holds her patient responsible for arriving at appointments sober. The philosophical literature that has addressed these cases has suggested that either we engage in pretense when we hold these agents responsible, or that they genuinely are responsible, but to a lesser degree than paradigm responsible agents.

I reject both of these descriptions of these cases. I argue that the pretense view doesn’t do justice to the phenomenology of holding marginal agents responsible. When parents hold their children responsible, most conceive of themselves as doing so genuinely, not merely with an eye towards behavior modification. Similarly, in therapeutic relationships, ideally therapists hold their patients responsible in a way that involves respect for their agency. The degrees of responsibility view also fails to account for an important feature of our responsibility practices with respect to marginal agents, namely that one’s relationship to the impaired agent has a significant impact on whether one should hold her responsible. While the parent should hold her five-year-old responsible for screaming, for example, the restaurant staff shouldn’t. My explanation of this is that relationships are sources of normative expectations about how both parties in the relationship should behave (i.e., relationship-based norms).

Relationship-based norms have two central components: their content and conditions of enforcement. Both components should be sensitive to agents’ capacities to understand and satisfy practical demands. Additionally, relationship-based norms derive their force from some token relationship, though their content needn’t be unique to that relationship. Indeed, relationship-based norms may have similar content in a lot of relationships of the same kind. This is due to the social significance of those relationships, which provides a set of narratives about what participating in a relationship of that kind involves (e.g., parents are generally tasked with disciplining their children in ways that others are not).

Though I’ve discussed childrearing here, relationship-based norms certainly aren’t just for kids. They have two general roles in our responsibility practices: First, they allow us to adjust expectations in relationships where at least one of the parties is a marginal agent or otherwise impaired. Second, some relationship-based norms concern matters that are strictly internal to a particular relationship (e.g., who takes out the trash, what fidelity requires in your romantic relationship), where creating these norms partially constitutes the relationship itself. Standard views of responsibility hold that agents are responsible solely in virtue of properties internal to them. My alternative picture claims instead that responsibility is a property we have in virtue of being in relationships with other people. Considering how our relationships with others animate our responsibility practices also allows us to appreciate a continuity between marginal and non-marginal agents as participants in the moral community.

Why Do We Care What Others Think? Should We?

James Swindler—Why Do We Care What Others Think? Should We?
Illinois State University

Our first task here is to get beneath the concepts of social status and social integrate that have animated analyses of sociality. Assuming that philosophical answers must explain why the phenomena in question are the case, we ask: why we care about what others think (of us) and whether we should. Caring is not just one attitude among others but the condition of having attitudes and caring what others think is the condition of having social attitudes, of social intentionality. We aim not merely to tell a story about how sociality develops or is maintained, but to say what it is.

Sociality entails caring how we seem and are to one another. In particular, we care about being worthy in others' eyes, about preserving dignity, our own and theirs. Meaningful social relations are enacted through social roles. Thus, sociality is a structure of roles among interdependent people and, derivatively, things. This gives credence to social integrate models. But again, what makes social roles, their incumbents, and the structures they constitute significant is the respect afforded them. This gives credence to status models. Clearly, one cannot effect a social role without being integrated into a social structure with the status accorded the role. Dignity attaches to roles, whether individual or group. Social roles are constructed for functional reasons, in pursuit of group ends. The dignity of role incumbents derives from free commitment to meaningfully contribute to realizing group ends. Incumbents enjoy powers and responsibilities not present absent sociality. Since every act enacts some role, free action inherently commits agents to role responsibilities, presuming the dignity of agency equally for all. Even simple language use divides us into speakers and hearers while affording dignity to both.

The dignity of roles is thus egalitarian. Our second task is to see that functional human sociality is thus deeply equivalent to morality. Agents who execute their roles solely for group reasons, indifferent to arbitrary moral distinctions, admitting therefore the moral equality of persons, act from what I call social intentions. A morally constituted group will comprise roles whose norms satisfy the demands of moral law to respect persons equally and incumbents whose intention is to satisfy the norms of their roles. Such a perspective is definitively both moral and social, demanding social intentions of each towards all, whether local or global. This means that there is a moral meta-role and, as human agents, we are all incumbents in that role. The description of that role is the moral law.

For confirmation, we ask what goes morally wrong in group failure and collectively wrong in moral failure. Examples from interstate highway traffic are considered phenomenologically. We conclude by questioning whether the idealism of morality exceeds that of sociality.

At the dawn of social philosophy, Plato raises our concern about what others think in Socrates' demonstration to Crito that a good man ignores the opinions of the self-serving *hoi polloi* who can act only by accident. It is no accident that we count immoral intentions as anti-social.

Judgment Aggregation and Collective Responsibility

Andras Sziget—*Judgment Aggregation and Collective Responsibility*
Linköping University & Lund University

Collectives are groups of human beings including formal organizations, various associations, and random assemblages of people. Responsibility-collectivism is the view that the moral responsibility of at least some such collectives is something over and above the aggregate moral responsibility of individual group members. According to responsibility-collectivists the following can be true: You have allocated all the moral responsibility there is to be allocated to all individuals involved, and still you have not allocated all the responsibility there is to be allocated. One influential argumentative strategy to show why such a “responsibility deficit” could arise is based on the analysis of cases which involve aggregations of individual judgments. It is now well-established that the aggregation of individual attitudes is subject to certain limitations. Impossibility theorems formally prove that there exists no aggregation procedure that could satisfy an intuitive set of criteria for deriving the collective view from those of the individual members. It is, however, not immediately obvious why these unavoidable limitations of judgment aggregation should be taken to support responsibility-collectivism. In fact, these impossibility theorems may be seen as adding grist to the responsibility-individualist’s mill since they seem to show that the collective’s view depends contingently on the procedure used to aggregate individual judgments. This finding might in turn be taken to strengthen the conclusion drawn from Arrow’s theorem and some interpretations of the the prisoner’s dilemma, namely that individually rational forms of behavior or attitudes can lose some of their rationally desirable properties when aggregated. Nevertheless, responsibility-collectivists have pressed into service these impossibility theorems in support of their position. They have advanced different arguments to forge a link between impossibility theorems and responsibility-collectivism. My first task in this paper is to show that there is actually a whole family of responsibility-collectivist arguments which are similar insofar as they take judgment aggregation as their point of departure but significantly differ in how they reach the collectivist conclusion. For example, some collectivists have argued that (i) the collective’s view (and resulting actions) can be morally objectionable, but (ii) given the potential disconnect between individual views and the collective’s view exposed by impossibility theorems, (iii) it is possible that only the collective as such can be held responsible for its view. Another argument holds that (i) given the possibility of such a disconnect, (ii) certain decision-procedures are to be favoured over others if the rationality of collective behaviour is to be ensured, and (iii) if such a decision-procedure is not chosen the collective will be responsible for failing to do so. Cataloguing these arguments is an important task because collectivists themselves do not discuss the differences among them. It is also important because of the second task I want to undertake in this paper which is to reject these arguments. I want to show that, once we tease them apart, they are either question-begging or they do not provide support for the responsibility-collectivist position. I conclude that impossibility theorems of judgment aggregation do not favour responsibility-collectivism.

Naturalness, Context and Social Categories

Elanor Taylor—Naturalness, Context and Social Categories

Johns Hopkins University

Certain features of social categories present challenges for attempts to formulate an ontology of social categories. In this paper, I describe three such features of social categories, and propose that we can accommodate them by understanding social categories in terms of context-dependent naturalness, an analogue of metaphysical naturalness that captures context-dependent, rather than metaphysical, structure.

Certain social categories are apparently explanatorily indispensable, in that we must recognize social categories if we are to understand certain social phenomena, including conflicts and instances of injustice. Some social categories are fragile, in that their existence is highly dependent upon particular features of a social context. If these highly specific contextual conditions change, then the category will disappear. For instance, social groups associated with particular trends are often very fragile. However, some social categories are less fragile, and these less fragile groups persist across national boundaries, different social classes, and across different historical eras. Social categories also appear to be metaphysically heterogeneous. There is an open, ongoing debate about realism about certain social categories, including gender categories. However, realism is far less plausible as an approach to tax bracket categories or beatniks. This indicates that a metaphysics of social categories might not treat all social categories as metaphysically alike.

These features generate puzzles. How can such highly fragile, context-dependent entities play such a significant role in social and historical explanation? How can we come up with a unified metaphysical approach to a class of entities that appear to vary so much? I propose that we can make sense of these features of social categories, among others, by appealing to context-dependent naturalness.

Metaphysical naturalness was famously defended by David Lewis. According to Lewis there is a spectrum among properties. At one end we have a privileged, special group that appear in the laws of nature, account for resemblance and causal powers and ground other properties, and at the other end are properties that are not like this. The former are the perfectly natural properties, and the latter highly non-natural properties. For example, on this view, the property of being negatively charged is more natural than the property of being to the left of the Eiffel tower on a Wednesday. Metaphysical naturalness is context-independent; if a property is perfectly natural it is so across all contexts. However, there is another phenomenon that looks very like metaphysical naturalness apart from the fact that it is context-dependent. In the same way that metaphysically natural properties are metaphysically privileged, certain properties are privileged with respect to a particular activity, but not necessarily in general. In my account of context-dependent naturalness I regiment this idea into a useful context-dependent alternative to metaphysical naturalness.

In this paper I describe and motivate the framework of context-dependent naturalness, and apply this framework to the case of social groups. I show that context-dependent naturalness can help us to make sense of certain intriguing features of social categories, while leaving room for further metaphysical enquiry about the nature of social categories. I also argue that context-dependent naturalness provides resources for understanding other features of social groups, including the nature of lawlike generalizations about social categories.

Group Knowledge Is Group Know How

*Deborah Tollefsen—Group Knowledge Is Group Know How
University of Memphis*

Who knows? Until recently, the answer to this question always referenced an individual agent. Sue knows, Abe knows, he knows, she knows. But developments in social epistemology, distributed cognition, and in the field of collective intentionality have opened the door to the possibility that groups can be epistemic agents and are the proper subject of knowledge ascriptions. In this paper, I argue that attempts to extend theories of individual propositional knowledge to groups are misguided. Group knowledge is best understood as the ability a group has to register that *p* accurately. In short, group knowledge is group know-how.

In section I, I begin by motivating the idea of group knowledge. In section II, I survey and offer a taxonomy of the types of theories of group knowledge on offer in the literature. In section III, I identify some challenges facing these approaches that have to do with the epistemic credentials of group belief. Given the attempt to extend theories of propositional knowledge to groups, these challenges are inevitable. Rather than attempt to stretch epistemological theories beyond recognition, I propose an entirely new path. In section IV, I explore this path by focusing on our practice of ascribing knowledge to groups. Identifying the contexts of group knowledge ascription will help to highlight the differences between human knowledge and group knowledge and will reveal that what we are really doing when we ascribe knowledge to a group is identifying an ability, an ability manifest through the various actions (joint and coordinated) of its individual members. A group's knowing that *p* is best understood as a form of know-how. I model this theory of group knowledge on a theory of individual knowledge developed and defended by Stephen Hetherington (2011). Hetherington argues for the controversial claim that all human knowledge—that is knowledge-how. Although controversial as an account of human propositional knowledge, it is a surprisingly useful and amenable theory of group propositional knowledge.

Silencing Indigenous Groups

*Leo Townsend and Dina Townsend—Silencing Indigenous Groups
University of Vienna and Tilburg University*

In a number of regional and domestic legal jurisdictions, courts and legislators have recognized an obligation on the state to consult with indigenous communities when making decisions or taking actions that will affect the community's land and/or environment. In international law, the right to consultation has been recognized as one that belongs to the community itself, rather than only to the individual members of the community. In other words, it is the community itself that is engaged and called upon to speak. But while such processes of consultation are intended to elicit and facilitate group speech, we argue that they often succeed only in silencing the groups in question. Not only are these groups seldom heard, they are often not even allowed to speak.

Our paper approaches this topic, of the impediments to effective community consultation, from both a philosophical and a legal perspective. We begin by outlining a broadly Austinian approach to group speech acts, and using it to distinguish three forms of group silencing. Locutionary group silencing occurs when a would-be group speaker is prevented from producing meaningful utterances in the name of the group; illocutionary group silencing occurs when a group's attempted speech act misfires because it is not recognized by the audience as an act of that kind; and perlocutionary group silencing occurs when the intended consequences of the group speaker's act are systematically forestalled. We then discuss three particular legal disputes, each one involving community consultation, in which these forms of silencing play a significant but under-appreciated role in disempowering the community's speech. First, we consider a dispute in which the community was a victim of locutionary silencing, since the consultation process sowed discord within the community and thus confounded any attempt on their part to come out with a genuinely collective statement. Next, we look at a case in which we find the community was silenced in the illocutionary sense, since their attempted assertions about the environment were taken instead as collective belief-reports, that is, claims about their cultural 'belief-system'. Finally, we examine a case in which the community was a victim of perlocutionary silencing, since although its assertions were correctly recognized as such, they were not taken even minimally seriously from an epistemic standpoint. We conclude the paper by arguing that while some efforts have been made to ensure communities are not the victims of locutionary silencing, more needs to be done to recognize and address the ways consultation allows for illocutionary and perlocutionary silencing of indigenous communities.

Can Social Groups Bear Obligations to Combat Structural Oppression?

William Tuckwell and Paul-Mikhail Catapang-Podosky—Can Social Groups Bear Obligations To Combat Structural Oppression?

The University of Melbourne

It is clear that groups that meet the conditions for being a collective agent can be the bearers of obligations. It is less clear that obligations can fall on highly unorganised groups, such as ‘the global collective’ (for an argument that suggests unorganised groups can be the bearers of obligations see Wringe 2014, and for an argument against see Lawford-Smith 2015). In between these two extremes lie many other groups of varying degrees of (un)organisation. Some of these groups are often said to be implicated in oppression, where oppression is to be understood as a form of injustice in which one social group is subordinated to another (Taylor 2016); for example, people of color are subordinated to white people and women are subordinated to men.

In this vein, last year former L’Oreal model Munroe Bergdorf declared that ‘all white people are racist’, but later qualified her claim by saying that ‘this isn’t about individuals’. Plausibly, the implication being that white people as a group, rather than individuals, are racist. Additionally, the Guardian journalist Jonathan Freedland has written that ‘if one thing is screamingly obvious here, it’s that men need to say to women that, when it comes to sexual harassment, the problem is not you – it’s us.’ Plausibly, the ‘us’ to which Freedland refers is men as a group rather than each individual man.

Drawing on Iris Marion-Young’s (2006) suggestion that participation in the processes that give rise to structural injustice grounds obligations to combat structural injustice, we aim to offer an account that makes sense of obligation attributions to uncoordinated social groups that are implicated in the creation and maintenance of oppressive social relations. We take the uncoordinated group ‘men’ as a test case and argue that ‘men’ bear obligations to combat the structural oppression of ‘women’ because, by their actions, ‘men’ contribute to the processes that give rise to, and constitute, structural oppression. We aim to show that the relevant processes include: reifying patterns of prejudicial credibility assessment (in Fricker 2007’s sense); upholding social meanings that structure social interactions in ways that advantage men and disadvantage women (Haslanger 2012); perpetuating asymmetric moral support roles (Manne 2017), and; enacting patriarchal social norms (Bicchieri 2016).

Constitutive Rules Are Not Reducible to Regulative Ones

Raimo Tuomela—Constitutive Rules Are Not Reducible to Regulative Ones
University of Helsinki

Hindriks and Guala (2015), pp. 472-474, claim that constitutive rules are non-creative and a kind of transformations of regulative rules. I claim that this is wrong. A plausible and intuitive view of a constitutive rule says that X constitutes Y and also what the Y-term in the Searlean formula "X counts as Y" refers to (e.g. that X thus is money). While a constitutive rule is definition-like, a regulative rule rather is supposed indicate how Y concretely will function and it thus normatively regulates people's behavior (cf. below).

The constitutive rule, C, in the Hindriks-Guala (H-G) 2015 paper it is claimed on p. 472 that if a person is the first to occupy a piece of land, it is his property and he has the exclusive right to use it. The regulative rule (R) involved here is regarded as a conjunction saying in effect that if one is the first to occupy a piece of land then, in a quasi-Lockean sense, it is constituted as his property and gives the owner the right to use it (see Searle, 2015, for criticism of this). Yet R does not say how the institution of property is to be used which is what a regulative rule should do (according to my account). As will be shown below, in the H-G account C and R on the basis of their sameness of form have basically the same truth-functional content. This implies that neither one is logically or conceptually primary.

Abbreviating the H-G regulative rule R as "If Occ, then Right" (p. 472) we get this: "If Occ, then Prop, and hence Right". The rule C has the form "If Occ, then Prop and if Prop, then Right". Combined the two conjuncts entail "If Occ, then Right", i.e. the rule R. R and C thus are equivalent.

Rules such as C and R qua, respectively, constitutive and regulative, are supposed to answer different questions and have different contents. However, the authors' above simple logical formalism in contrast entails that they have the same content. The rule C is indeed constitutive but in their account the regulative rule R is also constitutive because having the right involves a right-involving status function entailing deontic power. Yet R does not show how to regulate the use of the institution of property. The distinction between constitutive and regulative rules is real and not eliminable. The distinction is central. It is needed for the analysis of social institutions and also for "acting as a group" (in the sense of the members' collectively constructing or constituting themselves as a social group).

What has above been said about the case of property applies *mutatis mutandis* to the institutions of money and marriage, etc.

The Explanatory Power of Natural Kinds

Tuomas Vesterinen—*The Explanatory Power of Natural Kinds*
University of Helsinki

In this paper, I argue that the boundaries of natural kinds in the special sciences can be drawn based on the explanatory power of the maintaining mechanisms. According to the minimal interpretation of scientific realism, natural kinds should offer reliable inductive inferences and explanations. Many times, however, there is not just one right way to classify reality, but different taxonomies of the phenomenon come with excluding explanatory powers. Based on the homeostatic property cluster theory of natural kinds (Boyd 1989) and contrastive-counterfactual theory of explanation (Woodward 2003), I argue that classifications in the special sciences should be based on their accuracy as well as their explanatory power.

The HPC-theory has become the received view of natural kinds in the special sciences. According to the view, natural kinds consist of property clusters maintained by homeostatic causal mechanisms. However, the HPC-view has been criticised for not offering non-conventional ways to draw the boundaries of mechanisms and the level of their abstraction (Craver 2009). I argue that a pluralistic view of natural kinds, which is based on the explanatory power of the homeostatic mechanisms, dissolves these problems. Classifications that have more explanatory power are based on more stable kinds and enable more inductive inferences. I demonstrate that the boundaries of natural kinds depend in part on what the classification is aimed at explaining. This in turn determines the core explanandum properties and the right level of explanation. The dependence between kinds and explanations is more important in the human sciences than other special sciences (e.g. biology) because individual humans can fall under various unexclusive kind concepts. Moreover, there is no encompassing explanatory paradigm to embed all the human kinds (unlike evolution in biology) (Dupré 2004). Instead, human kind classifications serve different explanatory purposes that vary from mental disorders to gender.

However, when different classifications target the same explanandum, they may have explanatory trade-offs. The contrastive-counterfactual theory of explanation enables comparative analysis of explanatory powers. According to the theory, explanatory power is measured by the explanation's ability to answer relevant counterfactual what-if-questions. Ylikoski and Kuorikoski (2012) have analysed explanatory virtues based on the theory, and three of them are important for natural kinds: non-sensitivity, contrast precision, and accuracy. I provide examples from psychiatric classifications to analyse explanatory virtues. For example, schizophrenia can be classified differently depending on the purpose of the explanation, whether it is clinical treatment or universal classification.

Understanding Collective Emotions

Lovey Vikram—Understanding Collective Emotions

Jawaharlal Nehru University

The present work is an attempt to understand collective emotions within the framework of analytic philosophy. Here, I shall attempt to investigate the conceptual possibility of collective emotions. I shall try to examine the nature and the content of collective emotions by looking into its various aspects. The task at hand shall be accomplished by examining the nature of togetherness involved in the phenomena of collective emotions. Some of the related questions that this work is going to address within its ambit are: what do we refer to by the phrase “collective emotions”? Do we refer to a phenomenon in which each person’s emotional state is put together to constitute what is called collective emotion? Or, is there some such mental state which goes beyond the individual members and could be attributed to the collective as a whole? Can we sensibly talk about the possibility of an irreducible collective emotion?

In the proposed work, I shall attempt to analyse some of the popular responses towards the phenomenon of collective emotions put forth by philosophers and social scientists especially keeping in mind the various theories of emotions. For instance, those upholding a cognitive explanation of collective emotions such as Margaret Gilbert, Bryce Huebner, and others claim that the essence of collective emotions could be captured through a cognitive perspective alone. Some of the issues that draw our attention when we try to interpret the genuineness of collective emotional states within the framework of cognitivism are: Can collectives be said to bear the capacity to be intentionally directed at something? How do groups relate to paradigmatic intentional states like beliefs and desires? An individual can have an emotional state directed at some object, but can a group ‘proper’ be emotionally directed to a particular object? Could there be certain elements or constituents in collective emotions which could not be captured by the intentionalistic perspective such as feelings? In other words, can a group be considered to feel something? And so on.

While explaining individual emotions, many philosophers are of the view that emotions cannot be captured by the intentional or the feeling aspect alone. In fact, for a fairer and comprehensive understanding of emotions we need to consider both the aspects together that is, the intentional structure and affective experience of an emotion are typically intertwined. The ability to ‘feel’ plays a major role in the emotional process and could not be completely ignored or brushed aside as mere contingent accompaniment of the intentional structure. However, if any such association is possible with respect to collective emotions as well, is an issue worth philosophical investigation. Can such a hybrid account explain collective emotions as well? What could be the possible problems of accepting such a hybrid account of collective emotions? Does not the inclusion of the affective elements into the intentional structure hamper the cohesiveness of the phenomenon of collective emotions? These are important issues and need thorough investigation. This work attempts to analyse them extensively.

Social Construction of Gender in the Theory of Gender Performativity

*Jaana Virta—Social Construction of Gender in the Theory of Gender Performativity
University of Tampere*

In her famous book *Gender Trouble* (1990) Judith Butler presents a theory of gender performativity. The theory states that gender is constructed via constant iteration of gender performatives. Its basic argument is that gender is something we do and not something we possess or express. The theory states that gender is product of human action, and thus it is socially constructed. If the theory is taken seriously as a metaphysical theory of gender, certain questions, such as what is the gender that the theory is concerned about, need to be studied.

According to Sally Haslanger “[s]omething is a social construction [...] in case it is an intended or unintended product of a social practice” (86). Ron Mallon states it would be more useful to phrase social construction claims in the form of “X socially constructs Y” than in the traditional form “Y is socially constructed” (chapter 1). The former form shows explicitly that a) there is some X that does the construction, b) there is some Y that is constructed, and c) there is the active verb “to socially construct” that can refer to different forms of action in different situations. It can refer e.g. to causal form of construction or to constitutive form of construction, a distinction introduced by Haslanger (87). These distinctions can be used as a method or as tools while analyzing a specific social constructionist theory. If we use the distinctions to analyze the theory of gender performativity, the questions that arise are: who does the construction of gender, what is the gender that is constructed, and what form of construction is the construction effecting?

I argue that Butler’s theory can be divided into two parts:

- a) social construction of gender as an outer feature of a person, and
- b) the social construction of the illusion of a gendered innate essence in a person.

According to the theory gender as an outer feature of a person (a) is a constitutive form of social construction that is constructed mainly by the person whose feature it is, and which they construct by constantly performing gender performatives. While the illusion of a gendered innate essence (b) is a causal form of social construction that is constructed mainly by impersonal agents (media, practices, discourses, etc.), and which is caused by them constantly repeating gender performatives.

It is important to notice that these two aspects of gender are not the only aspects of gender there are. This means that the theory of gender performativity is not a theory about for example gender identity or of anatomical differences between sexes, and thus does not take a stand on their metaphysical nature.

Towards an Ontology of the Echo Chamber

*Konrad Werner and Krystian Wiciarz—Towards an Ontology of the Echo Chamber
University of Warsaw, Institute of Philosophy and Central Forensic Laboratory of the Police*

The aim of the talk is to provide an ontological account of a peculiar social construct known as the filter bubble, echo chamber, or information cocoon. In short, being in an information cocoon means being surrounded only by things, events, and data that fit one's expectations, beliefs, values, etc. In these circumstances, it is expected that one will be less willing to confront one's beliefs with those of others and view them as the only correct ones (see e.g. Zuiderveen Borgesius et al 2016; Matz et al 2017).

Information cocoons are genuine social constructs, and as such they deserve serious ontological treatment. Meanwhile, philosophical literature on the subject is scarce, with some noteworthy exceptions which, however, focus more on epistemological matters (e.g., Miller & Record 2013; Thi Nguyen 2018). Our aim is to take a step forward in filling this gap.

We start from the observation that there is a striking affinity between the basic structure of information cocoons or echo chambers on the one hand, and a phenomenon that has been known to biologists and philosophers of biology since the early 20th century, called "Umwelt" (Uexküll 1926), recently referred to as the "cognitive niche" (Clark 2006; Bertolotti & Magnani 2017; Werner 2017; see also Laland & Starelny 2006). In short, the cognitive niche of species X can be thought of as a realm composed of entities that are accessible to members of X because of their cognitive machinery. From a different angle, this is the task-domain in which members of X operate.

The proposed account draws primarily on the formal ontology of the niche set forth by Smith & Varzi (1999) as part of a broader theory of location, parthood and boundaries (Casati & Varzi 1999; Smith & Varzi 2000). Within this framework, we target the information cocoon as a certain pathology of the cognitive niche (the most recent proposal of Arfini, Bertolotti & Magnani 2017 goes in a similar direction); that is to say, a niche whose proper task of providing selective access to the surrounding world or data is replaced by selective blocking of access. The difference is subtle, yet pregnant: it is not about selection itself that we should be worried, since it is common to all kinds of cognitive pursuits, but about shared (or social) blocking. We argue that while teaming up in order to construct shared cognitive niches is propelled by the pursuit of getting more (and more diverse) data, information cocoons are peculiar cases in which a group's effort is dedicated to getting less (and less diverse) data. This mechanism is not specific to the age of the Internet, but it has recently been aided dramatically by technology. The aim of the presentation is to provide an ontological framework enabling us to understand this phenomenon.

Identity, Identification, and Intention in Formal Organizations

Randall Westgren—*Identity, Identification, and Intention in Formal Organizations*
University of Missouri

In organization theory, there is a growing reliance on identity and identification as explanans for collective behaviors of work teams and formal units of large organizations, and for the organization as a whole. One speaks of the "identity of" and "identification with" the formal organization. Organizational identity has its roots in the construct of social identity (Tajfel and Turner, 1979), where one's individual (personal or self-) identity is measured against social identities of groups to which the individual belongs. Sociologists are interested in such group memberships as gender, ethnicity, and occupation. Organization theorists are interested in more purposive memberships; i.e. to what extent is the self a function of the workplace, church, professional societies, cultural and social groups, and shared activities to which one purposely adheres. Organization theorists with training in social psychology and sociology have written extensively on the identity of firms as they relate to the employees, managers, and customers that comprise the formal organization and the measurement of coherence between self-identity of these members and the identity of the collective.

In the development of this literature, questions arise about whether the social (collective) identity is purposively constructed; arises organically with shared beliefs, goals, and ethos of the members; or is a metaphorical construct because the individual self cannot be shared in any meaningful collective manner.

Philosophers concerned with collective intentionality and group agency will recognize the disagreement over the reducibility of group features to individuals; at the center of this is the ontological status of we-intentions vs. I-intentions. It is easy to propose that individuals can identify with groups, without having to parse out whether the collective attributes and agency are merely summative or metaphorical. There are three recent developments in organizational identity research that can inform the scholarship on collective intentionality and agency: the social actor perspective, relational identity, and role identity. Each is distinct from traditional in-group/out-group social identity. The social actor perspective (cf. King, Felin, and Whetten, 2007) requires that the firm has identity and intentionality distinct from those of its constituent members. This is consistent with the List and Petit (2011) account of group agency, with significant opportunity for moving both accounts forward. Relational identity (Sluss and Ashforth, 2007) and role identity (Stets and Burke, 2000) are studied at the level of dyads and small groups, where interactions between individuals serving roles or positions drive shared intentions and collective actions. These models from social psychology can inform both the paradigmatic small-group accounts of Bratman and Tuomela. enhancing the understanding of their modest sociality and shared ethos, respectively.

The paper pushes the usefulness of role identity, relational identity, and the social actor identity modeling formal group behaviors rooted in social ontology. It takes group intentionality beyond team reasoning (Bacharach, Sugden and Gold). This paper is limited to formal organizations such as firms and avoids collective intentionality and agency around moving coffee tables, going to the cinema, and similar phenomena.

Voluntarism About Gender

Charlotte Witt—Voluntarism About Gender
University of New Hampshire

I have three goals. The first is to map out some of the terrain of what is at issue between voluntarists and ascriptivists about gender in relation to the political activities and theoretical demands of contemporary trans and feminist theorists and activists. The second is to provide a historical context for the contemporary conversation. And finally, I argue for the plausibility of an ascriptive account of social normativity in the case of gender.

Voluntarism about gender is on the upswing in feminist theory in the United States largely due to the important critical work of trans and feminist activists and theorists. For example, according to the Human Rights Campaign, gender identity refers to: “One’s innermost concept of self as male, female, a blend of both or neither – how individuals perceive themselves and what they call themselves. One’s gender identity can be the same or different from their sex assigned at birth”.

Although not stated in these precise terms the HRC statement amounts to a voluntarist view because it considers gender identity to be a sufficient condition for belonging to a gender: To identify as a woman is to be a woman. Notice that the HRC refers to one’s innermost concept of self and how one perceives oneself. Along the same lines, feminist and trans philosopher Talia Bettcher suggests the criterion of “sincere avowal.”

But what exactly is voluntarism concerning gender? Is it a view about the social construction of gender and the complex process (or processes) by which we become gendered? Or is it a perspective on the criteria for occupying a particular gendered social position or belonging to a gender kind? Or is it rather a view about our gendered social roles and why we are responsive to and evaluable under gendered norms? There is a proliferation of topics and perspectives that could come under the label of voluntarism about gender and many approaches to exploring the topic. Where to begin?

I think it is useful to situate contemporary discussions of gender voluntarism in relation to a rich and important historical resource. Gender voluntarism is not an idea from nowhere, but one with deep roots in feminist theory. So I begin with an extremely useful and insightful pair of essays by Judith Butler that offer an interpretation of Simone de Beauvoir. This will allow me to distinguish two types of gender voluntarism. I call them Kantian voluntarism and existential voluntarism. I then place recent theoretical work by Talia Bettcher on gender in the tradition of existential voluntarism. And, I conclude with an argument for the plausibility of an ascriptive account of social normativity in the case of gender.

What Can We Do? Collective Ability, Collective Obligation and Non-Agent Groups

Bill Wringe—What Can We Do? Collective Ability, Collective Obligation and Non-Agent Groups

Bilkent University

Some, though apparently not all, accept - correctly, in my view - that agential possibility is a pre-condition for moral obligation: that is that, in Kant's words, 'ought implies can'. Some also accept - incorrectly in my view - that sentences attributing agential capacities to groups can only be true when the groups in question are themselves agents, and infer that the only groups on which moral obligations can fall are agential groups. In this paper I'll explain how it can be correct to attribute agential powers to non-agent groups. A key claim will be that non-distributive plural attributions of agential capacities can often be understood as predications attributing capacities to pluralities of agents, and not to one collective agent. I'll call this the 'pluralist' (as opposed to the 'collectivist') conception of such capacity attributions.

I shall start by arguing that the 'pluralist' conception of certain sentences, such as 'Corbyn and May can make it a close election' is preferable to the 'collectivist' one. I'll then put forward a minimalist account of the conditions under which the pluralist reading is correct. I shall also argue - following a suggestion put forward in unpublished work by Anne Schwenkenbecher - that some of our intuitions about the capacities of non-agent groups are best accommodated not by raising the bar for a collective to have an agential capacity, but by regarding collective agential capacities as a matter of degree.

Finally, I shall explore how this line of thought bears on the possibility of whether non-agent groups, including what I have elsewhere called the 'global collective' can have obligations? That will depend, I shall suggest, on the kinds of considerations we take to underpin the 'ought implies can' principle. We take it to hold because of considerations about the relationship between obligation and blame. On this view incapacity will exculpate.

But we might also think that possession of capacity to a very low degree might also exculpate. This will count against the possibility of large non-agent groups, such as the global collective, having significant collective obligations. An alternative view, on which the 'ought implies can' principle derives from structural requirements on practical reasoning is, however, more plausibly compatible with the existence of obligations falling on the global collective.

“Racism Without Races”—Socio-Cultural Racialization, Latinxs, and Muslims

Daniel Brinkerhoff Young—“Racism Without Races”—Socio-Cultural Racialization, Latinxs, and Muslims

Department of Philosophy, New York University

This paper focuses on the claim that groups like Muslims and Latinxs are racialized in many contemporary societies. It is the claim that these groups are now treated as racial groups, or are like racial groups—but in what sense? I argue that Islamophobia and anti-Latinx animus are now predominantly forms of racialization in the sense that they involve the social construction of a new human kind category, where it is imagined that anyone belonging to that kind does so in virtue of sharing a relatively immutable property or set of properties passed on through relatively endogamous family structures. These forms of racialization, like the classic ones, amount to a mystifying and false ideology of racial essences. But I argue that, in contrast to pseudo-biological theories of race, these forms of racialization are socio-cultural—they imagine that the racial kind is grounded in a set of social or cultural facts passed on through families, rather than in biological ones. I therefore situate my argument within the current forms of social constructionism in philosophy but aim to draw attention to the ways in which socio-cultural racialization represents a related, though distinctive form of racial ideology that has received much less attention in the philosophical literature (though somewhat more in the social science literature).

The significance of this is two-fold. First, socio-cultural racialization represents the possibility of a social process very much like racism, but that does not rely on any belief in the 19th and 20th century notion of a biological race (hence a “racism without races”). This allows it much wider justificatory force in our present world in which this biological notion of races is at least officially, in most social and political institutions, censured. Although it is not by any means a new phenomenon, it seems to be particularly potent in the 21st century. Second, and relatedly, this allows socio-cultural racialization to replace the formerly pseudo-biological justification for much racist ideology beyond Islamophobia and anti-Latinx animus, as in the use of the “culture of poverty” thesis applied to African-Americans, for example.

This paper has three concrete aims. First, I aim to disambiguate the notion of racialization by providing a typology of racialization. Second, I aim to argue that the notion of one type of racialization in particular, which I call socio-cultural racialization, is needed to explain the social position of Latinxs and Muslims in many contemporary societies. I give some reasons for thinking that other types of ideological animus, like cultural intolerance, ethnocentrism, nativism, and xenophobia, are insufficient for explaining Islamophobia and anti-Latinx animus, and that we therefore require the notion of socio-cultural racialization.

Lastly, I argue that the mystifying character of socio-cultural racialization is less visible because the influential critique of social constructionists often targeted notions of natural essences in favor of explanation based on social facts, but socio-cultural racialization tends to glide under the radar by basing its claims on the very same types of facts—social ones.

Why Conventions Cannot Be Social Institutions?

Vojtěch Zachník—*Why Conventions Cannot Be Social Institutions?*

University of Ostrava

Conventions, sometimes described as descriptive norms (Bicchieri, 2005; Ullmann-Margalit, 1977), play an important role in social interactions. They help people to coordinate their actions and solve equilibrium selection problem (famously Lewis, 1969 or Sugden, 1986), however, given the recent interest in social ontology (Guala, 2016; Epstein, 2015), there is a general issue whether conventions can belong to a category of social institutions. If we start thinking about conventions as a one member of broad set containing cases such as money, marriage, social norms, governments, etc., it may be evident that some similarities occur with regard to structure, function, and obligations.

Despite the mentioned variety of institutional facts there is a need to address the crucial issue of how to clarify both: the notion of convention and the notion of social institution. Only more lucid idea may help philosophers and social scientists to analyse relation of these social phenomena with precision. At first, I introduce five distinctive features of conventionality that delimit the notion and provide some understanding without any link to a specific definition; namely: interdependence, arbitrariness, independence on mental representations, normative neutrality, and spontaneity. Then, I will present the two major and the most prominent accounts of social institutions that explain them in terms of rules (North, 1990; Searle, 1995), or equilibria (Sugden, 1986; Binmore, 2010). My general strategy is the following: I will hypothetically assume that each account of institutions is true in order to confront this finding with beforementioned characteristics of conventionality. Consequences drawn from this procedure will shed a light on compatibility of the two social categories.

In a nutshell, I argue that social conventions are neither rules, nor equilibria, which raises a concern whether they can be institutions. And it seems unlikely that they can be institutions as described by these two prominent accounts, instead they share some features with both of them and provide a middle ground (similarly to Greif, Kingston, 2011). Therefore, I intend to emphasize the fact conventions cover a broad area of social behaviour that results from gradual process of social interactions as well as from “rules-like actions”. Under these circumstances, I deem unfruitful to claim that conventions are institutions, even though some cases of conventionality could perfectly fit into one of the two theoretical standpoints on the nature of institutions.

Overall, this result should not be considered as an unsatisfactory result since it allows us to rethink a role conventions play in social world. Many have struggled with this issue as they thought that conventions are something primitive and basic that helps to build more complex social constructions, or on the other hand that conventions are more advanced and similar to social norms. I intend to show that conventionality is multilayered phenomenon that reveals a novel and significant way how to approach social ontology.

Conceptual Analysis and the Nature of Law

Aleardo Zanghellini—Conceptual Analysis and the Nature of Law
University of Reading

This paper uses central case methodology and the philosophy of language to illuminate the role of conceptual analysis in law. It argues that the nature of law (the kind of thing that law is) and the nature of conceptual analysis (the kind of method that conceptual analysis, as currently practiced by legal philosophers, is) make something like the latter indispensable to systematic knowledge about the former. Or, at least, this must remain so barring a radical paradigm shift in our understandings about the constructedness of the social world.

Despite Leiter's and others' calls for 'naturalizing' jurisprudence, conceptual analysis remains, and is likely to remain, the methodology of much contemporary mainstream (non-normative) jurisprudence. Two reasons combine to account for this. The first has to do with the nature of law and legal things. The second has to do with the way in which conceptual analysis has evolved in the second half of the twentieth century, its function no longer being limited to yielding purely analytic truths, that is, truths by virtue of meaning alone.

I elaborate on these points by drawing on classical philosophy of language and on central case methodology. Classical philosophy of language, with its distinction between the sense and referent of a term, and between conventional kinds and natural kinds, suggests that some entities are more highly 'constructed' than others, and that law belongs to the former kind. Central case methodology, on the other hand, helps us see that:

- a) there are central cases and non-paradigmatic cases of conceptual analysis, depending on whether the legitimate province of conceptual analysis is seen as being limited to yielding analytic truths or as extending beyond that; and
- b) ever since Hart, conceptual analysis in law has departed from the central case: it has purported to yield synthetic, and not purely analytic, truths.

I conclude by arguing that the nature of law is such that:

- a) some aspects of law are impervious to conceptual analysis, but not to empirical social science; but also
- b) some aspects of law are impervious to empirical social science, but not to conceptual analysis.

It follows that conceptual analysis may not be sufficient, but is necessary, to a systematic understanding of law.

Which Room Am I In? Arendt and Husserl On Collective Contexts and Homeless Minds

David Zoller—*Which Room Am I In? Arendt and Husserl On Collective Contexts and Homeless Minds*

California Polytechnic State University, San Luis Obispo

Hannah Arendt offers in the case of Adolf Eichmann a study in what she views as banal, and culpable, thoughtlessness. Eichmann's preoccupation with mundane office politics leads him to neglect, and disclaim responsibility for, the endpoint of his group's actions: the deportation of millions of people to Nazi death camps. Eichmann's case indicates a puzzle behind current conversations on collective action and responsibility. It is commonly thought—though recently questioned by George Sher—that the "searchlight" of one's perceptual field has purchase on what is responsible for. Sher-style counterexamples offer cases where one is culpable for an artificially restricted "searchlight": for instance, the thoughtless individual leaves her dog in a hot car, but is intuitively responsible for what she didn't notice. If Sher's examples—and commonplace conclusions that hold "Eichmann's" culpable for collective outcomes—offer easy critiques of the "searchlight" model, Arendt's model of the philosopher tests how far we can go in that direction. Arendt's image of the "cosmopolitan," "homeless" self of the thoughtful philosopher depicts a mind literally everywhere and nowhere at once: nowhere because her immediate "searchlight" has no bearing on what she takes herself to be concerned with and responsible for; everywhere, in that everything is relevant.

There are clear reasons for discarding the "searchlight" view. To supplement Sher's reasoning, the confines of the searchlight likely contain some crucial moral information only in an attenuated form that strikes the agent as irrelevant compared to whatever more "real," immediate situation she feels she has to deal with. Eichmann's case suggests that that constriction of vision is unacceptable. On the other hand, discarding the searchlight as a guide to my responsibilities in collective settings renders me "homeless," in Arendt's terms, with no means of determining "what room I am in" morally speaking. Without the "searchlight," my personal moral situation in collective settings is indeterminate. While we could argue that such a "homeless" mental life is unpalatable, the unpalatability of the alternative cannot justify the searchlight view. Neither path seems acceptable.

I suggest a possible, controversial, resolution in Husserl's ethics of vocation and sacrifice: not all goods are compossible, and every good is sought at the expense of another. My situation is not universal, and if we have a responsibility to "think" widely about our commitments—as Eichmann does not—we nevertheless have reason to think in light of commitments to specific projects and groups. Since there is no way to "get things right" in any case, we must accept a tragic element in working from our perceptions and particular commitments.