

1st European Experimental Philosophy Conference

XPhi Europe 2021

17th – 19th June 2021, online

Book of Abstracts

<https://xhiprague.eu/>

Contents

A Pan-Cultural Distaste for the Wisdom of Crowds	4
Affect-driven ascription of <i>mens rea</i> : legal experts and laypeople surveyed	4
An Experimental Study on Perspectives in Equity Theory.....	5
Are intuitions about lying prone to partisan bias? An empirical investigation of the role of political beliefs in people’s judgments of lying	5
Belief in mind-body dualism and other minds	6
Can Artificial Intelligence Make Art?	6
Causal judgment in the wild: evidence from the 2020 US presidential election	7
Causation and Norms	7
Causation and Responsibility – A New Entailment Relation	8
Cognitively Modelling Transformative Experiences	8
Context-Sensitivity of Color Adjectives and Folk Intuitions	8
Counterfactual Thinking and Recency Effects in Causal Judgment.....	9
Degrading Causation	10
Demanding the Morally Demanding: Gender Differences in Response to Morally Demanding Charitable Solicitations.....	10
Disagreement is Said in Many Ways	11
Do Thick Concepts Provide Reasons for Action?	11
Does Macbeth See a Dagger? An Empirical Argument for the Existence-Neutrality of Seeing	12
Does Nobody Know What a Moral Judgment is? Education, Cognitive Reflection, and the Moral-Conventional Distinction	12
Epistemic mental states elucidated through electrophysiological responses. An experimental approach to Philosophy of Mind	13
Epistemic Side-Effect Effect: A Meta-Analysis.....	13
Examining Evaluativity in Legal Discourse: A Comparative Corpus-Linguistic Study of Thick Concepts	15
Fragmented and Conflicted? Folk Beliefs about Vision.....	15
Gratuitous Risk: perceived danger and recklessness judgements about outdoor sports participants	16
Guilty Artificial Minds	16
How Do People Balance Death against Lesser Burdens?	17
Law, Coercion, and Folk Intuitions	18
Morality Rules: Rule Application is Shaped by Moral Evaluation	18
Outcome Effects, Moral Luck and the Hindsight Bias	19
Outcome-valence and the reasonable person standard.....	19
Philosophical Expertise Put to the Test	19
Psychological Implications of Gendered AIs and Their Social Impact	20
'Really Know': A Linguistic and Philosophical Analysis	20

Revisionist attitude reports and inference patterns	21
Should I say that? An experimental investigation of the norm of assertion.....	21
Socratic Questionnaires.....	22
The Asian Disease Problem and Doing/Allowing Harm.....	22
The Future of Virtual Reality in Experimental Research: Revolutionary or Problematic?	23
Theoretical virtues: do scientists think what philosophers think they ought to think?.....	23
Tracing Thick Concepts Through Corpora	24
Weakness of Will: Existence and Rationality.....	24
Western Skeptic vs Indian Realist. Cross-Cultural Differences in the Zebra Case intuitions	25
What are the components of a meaningful life?	25
What is the folk concept of life?	26
When is reappropriation appropriate? Lessons from 'gossip'	26

A Pan-Cultural Distaste for the Wisdom of Crowds

Minha Lee, Seoul National University, lee.minha@gmail.com

Anna Dorfman, University of Waterloo, anna.dorfman@uwaterloo.ca

Kelli Barr, University of Pittsburgh, kelli.barr@pitt.edu

Kristopher Smith, University of Pennsylvania

Steve Stich, Rutgers University

Clark Barrett, UCLA

Igor Grossmann, University of Waterloo, igrossma@uwaterloo.ca, www.igorGrossmann.com

Edouard Machery, University of Pittsburgh, machery@pitt.edu, www.edouardmachery.com

Under many conditions, aggregating opinions increases the accuracy of judgment. What's more, in at least some circumstances people are inclined to take others' judgments and decisions into account when they themselves have to make a judgment or a decision. But how do people view crowdsourcing—the practice of soliciting the opinions or preferences of many (often anonymous) others to inform one's judgments and decisions? Do people's folk theory about how one should make decisions reflect both the normative significance of crowdsourcing and people's inclination to take others' advice into account? And do these folk theories vary across cultures with an independent vs. interdependent social orientation and in small-scale vs. large-scale societies? To address these questions, we sampled 2,763 student and non-student participants in 11 industrialized and small-scale societies. It turns out that, all over the world, the folk theory of decision making is hostile to crowdsourcing, despite both the effectiveness of aggregating opinions and people's demonstrated conformism. We also document a self-other bias in this folk theory: people view decisions based on personal (rather than others') opinions as wiser. In sum, folk theories of everyday decision-making are diametrically opposed to strategies that actually most likely to promote effective decision-making.

Affect-driven ascription of *mens rea*: legal experts and laypeople surveyed

Marc-André Zehnder, University of Zurich, Center of Ethics UZH Philosophisches Seminar, Zollikerstrasse 117, CH-8008 Zürich, marcandre.zehnder@uzh.ch, https://www.philosophie.uzh.ch/de/seminar/people/research/ethics_schaber/zehnder.html

An affect-driven information processing might explain why professional judges are affected by the outcome severity in their ascription of intention (Kneer & Bourgeois-Gironde, 2017), although culpable mental states (*mentes reae*) should be judged independently from the forbidden act (*actus reus*) and its outcome. Similar processes could occur if judges or juries get confronted with other affect-laden information, like the suspect's moral character, which should neither be considered in ascribing *mens rea* (cf. e.g., Nadler, 2012).

I will present two studies, one with laypeople and one with legal experts, reporting the first empirical investigation into the affective-mediated ascription of *mens rea* (recklessness and negligence) made by legal experts.

We found that laypeople and legal experts reported stronger negative affect, ascribed more blame to the suspect, and were more willing to ascribe inculpatory states of mind if the suspect had a

bad moral character than if the suspect had a good moral character. Interestingly, after evaluating further evidence, both moral character and outcome severity had bigger effects on the ascription of blame and *mens rea* in the study with experts than in the study with laypeople, although experts report lower levels of initial blame and negative affect. Possible explanations will be discussed.

An Experimental Study on Perspectives in Equity Theory

Kazutaka Inamura, School of Political Science and Economics, Waseda University, kinamura@waseda.jp

Michiko Ueda, School of Political Science and Economics, Waseda University, mueda@waseda.jp

Robert F. Veszteg, School of Political Science and Economics, Waseda University, rveszteg@waseda.jp

Keywords: equity theory; desert; justice; impartiality; egalitarianism; redistribution

This article reports statistical results from a representative survey experiment with 1900 respondents in Japan about how the first- and third-person perspectives affect people's normative choices regarding a fair distribution of scarce resources. Our goal is to perform an empirical test of equity theory (or desert-based theory of justice) according to which "inputs (into a relationship) and outputs (out of a relationship) should be in the same proportion for all persons involved" (Gaertner and Schokkaert 2012, p.98). We have created 18 vignettes by altering the values of three treatment variables: perspective (first vs third person), relative position (advantaged vs disadvantaged), and context (distributing bananas on an uninhabited island, assigning bonus payment in a business, and redistributing labor income through taxation). We find that decision-makers in the first-person perspective tend to be more egalitarian, willing to sacrifice personal gain for helping the other in a disadvantaged position. We do not find a similar effect in the third-person perspective. Regression analysis reveals that socio-demographic variables do not have significant power in explaining the observed outcomes. We conclude that, for people's normative choices regarding a fair distribution, perspective does not matter, but relative positions do.

Are intuitions about lying prone to partisan bias? An empirical investigation of the role of political beliefs in people's judgments of lying

Louisa Reins, Ruhr University Bochum, louisareins@icloud.com

Lying is a familiar and important moral phenomenon that most people are affected by on a daily basis—be it in personal relationships, election battles, lawsuits, or in the form of fake news. In the present paper, we examine whether people's judgments of lying are prone to partisan bias, which describes the tendency of people to evaluate identical information more favorably when it supports rather than challenges their political allegiances and beliefs. In two preregistered experiments, we investigate the hypotheses that participants give harsher lie judgments when deceptive statements are uttered by political opponents (Experiment 1) and/or express an opposed political belief (Experiment 2), both in straight-out lies and deceptive implicatures. The latter provide a particularly interesting context to investigate partisan bias, since they are commonly used in political discourse and leave room for deliberation about whether or not they should be counted as an instance of lying.

The practical and philosophical implications of our findings (e.g., with regard to the question of whether folk intuitions about lying are reliable) will be discussed.

Belief in mind-body dualism and other minds

Pascal Burgmer, University of Kent, United Kingdom, p.burgmer@kent.ac.uk,
<https://www.kent.ac.uk/psychology/people/1434/burgmer-pascal>

Matthias Forstmann

Andrew Todd

People differ in their beliefs about how minds relate to bodies, and this variability in lay conceptions of mind-body relations has important psychological consequences. Drawing on the proposition that a dualistic view on mind-body relations is a developmental by-product of learning to reason about other minds, we report five studies (N = 1,459) investigating the bi-directional relationship between belief in dualism and thinking about other minds. Participants who endorsed dualism viewed the mind as an entity that is primarily accessible through mentalizing, but not through the observation of bodies (Studies 1 & 2). Conversely, actively mentalizing (thinking about the mental states of another person; Study 3) and empathizing (feeling what another person might be experiencing; Studies 4 & 5) facilitated belief in mind-body dualism. Finally, empathizing selectively enhanced belief in substance dualism (vs. reductive physicalism), which aligns with the notion that minds are immaterial entities that can exist in the absence of physical bodies (Studies 4 & 5). Collectively, these findings suggest an intimate relationship between viewing unobservable minds as separate from observable bodies and reasoning about the minds of others.

Can Artificial Intelligence Make Art?

Elzė Sigutė Mikalonytė, Institute of Philosophy, Vilnius University, elze.mikalonyte@fsf.vu.lt

We ran two experiments (total N=693) to explore whether people are willing to consider paintings made by AI-driven robots as *art*, and whether they are willing to consider robots *artists*. We manipulated three factors: (i) agent type (robot v. human agent) so as to contrast the robot results with those for ordinary human artistic creation; (ii) behavior type (intentional creation of a painting v. accidental creation), so as to explore the role intentions and mental capacities might play in the ascription of artistic value and agency, and (iii) object type (abstract v. representational painting), which might interact interestingly with behavior type. Besides artistic merit of the object and artistic agency of its creator, we also measured whether people were as willing to ascribe certain relevant mental states (belief, desire, intention) to the robot as to the human being.

According to our main findings, people are by and large as willing to consider robot paintings as art as human paintings. They are, however, much less disposed to consider robots as artists than humans. Our further results suggest that artistic intentions are a necessary criterion for artistic agency, though for an object to be art, such intentions play a less important role.

Causal judgment in the wild: evidence from the 2020 US presidential election

Tadeg Quillien, School of Informatics, University of Edinburgh, tadeg.quillien@gmail.com,
<https://sites.google.com/view/tadeg-quillien>

Michael Barlev

Preprint: <https://psyarxiv.com/7w9re/>

A given event has many causes, but people intuitively view some causes as more important than others. Models of causal judgment have been evaluated in controlled laboratory experiments, but they have yet to be tested in complex real-world settings. Here, we provide such a test, in the context of the 2020 U.S. presidential election. Across tens of thousands of simulations of possible election outcomes, we computed, for each state, an adjusted measure of the correlation between a Biden victory in that state and a Biden election victory. These effect size measures accurately predicted the extent to which U.S. participants (N=207, pre-registered) viewed victory in a given state as having caused Biden to win the presidency. This supports the theory that people intuitively select as causes of an outcome the factors with the largest average causal effect on that outcome across possible counterfactual worlds.

Causation and Norms

Levin Güver, Faculty of Law & Department of Philosophy, University of Zurich, Switzerland,
levin.guever@hotmail.de

Markus Kneer, Department of Philosophy, University of Zurich, Switzerland,
markus.kneer@gmail.com

www.guiltymindslab.com

When two agents jointly bring about a bad outcome, ordinary folk will consider the norm-violating agent as more causal and label her *the* cause of the outcome. Two camps purport to explain this effect. The *responsibility* view (RV) holds that the ordinary concept of causation is normatively valenced. The *bias* view (BV), on the other hand, maintains that said concept is descriptive, and that the folk are falling victim to a pervasive blame-driven bias. Both the RV and BV make identical predictions for the vast majority of cases—only the BV, however, predicts that features *peripheral* to agential moral responsibility will have an influence on folk causal judgement. In our paper, we set out to test a novel peripheral feature: That of *norm-type*. In a series of experiments (N=600), we show that ordinary causal judgement is, against predictions made by the RV, sensitive to the *type* of norm violated—be it a pertinent, non-pertinent, or wholly silly norm—and that this effect is best explained by an evoked desire to blame the agent. The practical relevance of this research culminates in the law, which imposes legal liability on grounds of what we take to be *biased* ascriptions of causality.

Causation and Responsibility – A New Entailment Relation

Justin Sytsma, Philosophy Programme, Victoria University of Wellington, justin.sytsma@vuw.ac.nz, <http://justinsytsma.com/>

Pascale Willemsen, Department of Philosophy, University of Zurich, pascale.willemsen@philos.uzh.ch, <https://www.pascalewillemsen.com/>

Kevin Reuter, Department of Philosophy, University of Zurich, kevin.reuter@uzh.ch, <http://www.kevinreuter.com/>

The standard view in philosophy is that responsibility entails causation. Most philosophers treat this entailment claim as an evident insight into the ordinary concept of responsibility. Further, it is taken to be equally obvious that the reversal of the entailment claim does not hold: causation does not entail responsibility. In contrast, the responsibility account put forward by Sytsma and Livengood predicts that the terms ‘responsible for’ and ‘caused’ will generally be taken to apply in the same contexts. If the responsibility account is correct, then the reversal of the entailment claim may hold, and, a fortiori, there would be mutual entailment between responsibility and causation. Using the cancellability test, we provide empirical evidence that causation and responsibility are mutually entailed by each other.

Cognitively Modelling Transformative Experiences

Marvin Mathony, Berlin School of Mind and Brain

Michael Messerli, University of Zurich

In the last six years, philosophers have discussed the topic of transformative experiences. In this paper, we contribute to a crucial issue that is currently under researched: the failure of cognitive modelling. We argue that cognitive modelling can be operationalized as affective forecasting, and we compare transformative and non-transformative experiences with respect to the ability of affective forecasting. In particular, our claim relies on a meta-study conducted by Levine et al. (2012) on affective forecasting. We looked into each of the 84 studies listed in their meta-study to categorize them based on the criterion of whether the affective forecasting task was concerned with transformative or non-transformative experiences. The result found in a two one-sided equivalence test (TOST) is that decision-makers’ performance in cognitively modelling transformative experiences does not systematically differ from decision-makers’ performance in cognitively modelling non-transformative experiences. This claim stands in strict opposition to L.A. Paul’s main argument.

Context-Sensitivity of Color Adjectives and Folk Intuitions

Adrian Ziółkowski, University of Warsaw, Faculty of Philosophy, Experimental Philosophy Laboratory *KogniLab*, adrian.a.ziolkowski@uw.edu.pl, www.kognilab.pl

keywords: contextualism – semantic intuitions – context-shifting experiments – color adjectives – context sensitivity

In this talk I report new empirical data on folk semantic intuitions concerning color adjectives in the so-called “context-shifting experiments”. Contextualist use context-shifting experiments to argue that context can shape meaning and truth-conditions to such a degree that might result in opposite truth-evaluations of the same sentence in different contexts. Initial findings by Hansen and Chemla (2013)

suggest that laypersons' semantic judgments are sensitive to context in the way predicted by contextualists. I focus on context-shifting experiments that involve color adjectives and present experiments that are a partial replication and methodological extension of Hansen and Chemla's study. One aim was to corroborate their findings using a bigger sample, but the main goal was to test the stability of results in different methodological variants of empirical adaptations of context-shifting experiments. As suggested by Ziółkowski (2017), certain experimental settings (within-subjects) might bring data more favorable to contextualism than other settings (between-subjects). My study compares three experimental settings: within-subjects (with randomized order of context presentation), between-subjects (different contexts are presented to distinct groups), and "contrastive design" (both contexts are presented side-by-side). My results are highly consistent across the methodological variants I employed, but while they show some effects expected by contextualists, it is disputable whether they bring strong support to contextualism.

Counterfactual Thinking and Recency Effects in Causal Judgment

Paul Henne, Lake Forest College, p_henne@lakeforest.edu, www.paulhenne.com

Preprint: <https://osf.io/vgbkx/>

People tend to judge more recent events, relative to earlier ones, as the cause of some particular outcome. This recency effect, however, reverses in cases of overdetermination: people judge that earlier events, rather than more recent ones, caused the outcome when the event is individually sufficient but not individually necessary for the outcome. In five experiments, we find evidence for the recency effect and the primacy effect for causal judgment. Traditionally, these effects have been a problem for counterfactual views of causal judgment. However, we argue that an extension of a recent counterfactual model of causal judgment explains both the recency and the primacy effect. In line with the predictions of our extended counterfactual model, we also find that, regardless of causal structure, people tend to imagine the counterfactual alternative to the more recent event rather than to the earlier one. Moreover, manipulating this tendency affects causal judgments in the ways predicted by this extended model: asking participants to imagine the counterfactual alternative to the earlier event weakens the interaction between recency and causal structure, and asking participants to imagine the counterfactual alternative to the more recent event strengthens the interaction between recency and causal structure.

Degrading Causation

Kevin O'Neill, Department of Psychology & Neuroscience, Duke University, kevin.oneill@duke.edu, kevingoneill.github.io

Paul Henne

Paul Bello

John Pearson

Felipe De Brigard

Preprint: <https://osf.io/m5w9c>

When asking if lightning caused the forest fire, one might think that the lightning is more of a cause than the dry climate (i.e., it is a graded cause) or they might instead think that the lightning strike completely caused the fire while the dry conditions did not cause it at all (i.e., it is a binary cause). Psychologists and philosophers have long debated whether such judgments are graded. To address this debate, we started by reanalyzing data from four recent studies. In this context, we provide novel evidence that causal judgments are actually multimodal: although most causal judgements were binary, there was also some gradation. We then tested two competing explanations for the gradation we observed: the confidence explanation, which states that gradation distinguishes between certain and uncertain causes, and the strength explanation, which states that gradation distinguishes between strong and weak causes. Experiment 1 tested the confidence explanation and showed that gradation in causal judgments was moderated by confidence. People tended to make graded causal judgments when they were less confident, but they tended to make discrete causal judgments when they were more confident. Experiment 2 tested the causal strength explanation and showed that although causal judgments varied with factors associated with causal strength, confidence ratings were unchanged. Overall, we found that causal judgments are multimodal and that observed gradation reflects independent effects of confidence and causal strength on causal judgments.

Demanding the Morally Demanding: Gender Differences in Response to Morally Demanding Charitable Solicitations

Philipp Schönegger, University of St Andrews, ps234@st-andrews.ac.uk, <https://philipp-schoenegger.weebly.com/>

Ben Grodeck, Monash University

Should we confront people with morally demanding statements to perform certain actions. Specifically, should we tell people that they have a moral obligation to donate to charity? To investigate this question, we conduct an online randomized experiment via Prolific (n=2500) where participants can donate to charity. Using a between-subject design, we provide some participants with the same moral argument as to why they should donate. We then add and vary the level of the moral demandingness to donate across treatments by adding a single sentence. We find that moral arguments increase the amount donated by 51.7%. Further, we find that increasing the level of the moral demandingness does not affect the frequency or amount of donations. However, the highest level of moral demandingness (“You are morally obligated to donate”) has disparate effects on men and women. We find that 14% more women make a donation when faced with this nudge of high moral demandingness while men donate 31.9% less. These findings suggest that demanding the

morally demanding may have disparate effects on different populations and may as such be applicable only in certain contexts.

Disagreement is Said in Many Ways

David Bordonaba-Plou, FONDECYT postdoctoral researcher, Universidad de Valparaíso, Chile, davidbordonaba@gmail.com, <https://uv-cl.academia.edu/DavidBordonabaPlou>

During the last years, one theme that has attracted the most attention in philosophy of language has been *taste disagreements* (see Kölbel, 2004; Lasersohn, 2005; MacFarlane, 2014), which have been depicted using examples like:

- 1) Laila: Artichokes are tasty.
- 2) David: No, artichokes are not tasty.

In most of these examples, the way the second speaker expresses disagreement is always an *explicit mark of disagreement*, simple negations like ‘No,’ or other expressions like ‘I disagree,’ or ‘That’s not true.’

This work defends that, although most authors have used this characterization, it is not adequate to represent taste disagreements because taste disagreements are expressed using a vast array of expressions, and not only the structure ‘negation (No/I disagree/That’s not true/That’s false) + x is not ϕ .’ I will build a corpus of taste discussions using Reddit threads; then, I will run several analyses to determine the diversity of ways used by speakers to express disagreement. Specifically, I will show, first, that the structure ‘negation + x is not ϕ ’ is not common in the corpus and, second, that speakers express disagreement in many different ways, e.g., using rhetorical questions, constructions with ‘miss,’ ‘prefer,’ or ‘wrong,’ or ironic statements.

Do Thick Concepts Provide Reasons for Action?

Judith Martens

Pascale Willemsen

In this talk, we will present three pre-registered studies on the use of thick concepts in ordinary language and on how they relate to reasons for actions. More specifically, we investigate whether statements containing thick terms provide reasons for action to change one's behaviour or stick to it. The first two studies examine whether thick concepts have the disposition to be interpreted as reasons for action. We presented participants with a short story about an agent who needs to make a decision. In Experiment 1, the agent herself thinks that performing a certain action would be, for instance, cruel, rude, honest, etc. in Experiment 2, we moved from a first-person to a third-person perspective in which a second person utters that performing a particular action would be cruel, etc. Both studies demonstrate that thick concepts are readily interpreted as providing reasons for actions. A statement containing a thick concept is interpreted as more action-guiding than any descriptive term that we tested and just as action-guiding as thin terms.

How are these reasons for actions communicated? There are at least two obvious options on the table. First, reasons for actions are communicated as a matter of the semantic meaning of thick

concepts. Alternatively, uttering a thick term might conversationally implicate reasons for actions. We examined these two options in a third experiment. Our results strongly suggest that the action-related, reason-giving component of thick concepts is only loosely connected to the thick concept. In addition, we also found a significant polarity effect, such that negative terms seem to be more action-guiding than positive terms.

Does Macbeth See a Dagger? An Empirical Argument for the Existence-Neutrality of Seeing

Vilius Dranseika, Jagiellonian University / Vilnius University, vilius.dranseika@uj.edu.pl, www.dranseika.lt

André Sant'Anna, Washington University in St. Louis, <http://andre-santanna.com>

Preprint: <https://philarchive.org/archive/SANDMS-3>

In a recent paper, Justin D'Ambrosio (2020) has offered an empirical argument in support of a negative solution to the puzzle of Macbeth's dagger—namely, the question of whether, in the famous scene from Shakespeare's play, Macbeth sees a dagger in front of him. D'Ambrosio's strategy consists in showing that "seeing" is not an existence-neutral verb; that is, that the way it is used in ordinary language is not neutral with respect to whether its complement exists. In this paper, we offer an empirical argument in favor of an existence-neutral reading of "seeing". In particular, we argue that existence-neutral readings are readily available to language users. We thus call into question D'Ambrosio's argument for the claim that Macbeth does not see a dagger. According to our positive solution, Macbeth sees a dagger, even though there is not a dagger in front of him.

Does Nobody Know What a Moral Judgment is? Education, Cognitive Reflection, and the Moral-Conventional Distinction

Ross Colebrook, Baruch College, City University of New York, ross.colebrook@baruch.cuny.edu, www.rosscolebrook.weebly.com

Paper:

<https://www.dropbox.com/s/6lxiq674oagxyy8/Does%20nobody%20know%20what%20a%20moral%20judgment%20is.docx?dl=0>

Using researcher-selected batteries of judgments, social domain theorists have revealed a wealth of evidence that people conceive of moral judgments differently from social-conventional judgments. But some moral psychologists, skeptical about the methods of social domain theory, have begun investigating moral judgments using a participant-driven paradigm, where researchers give a battery of judgments to participants, who then classify those judgments (as moral, conventional, personal, etc.) for themselves. This paradigm reveals a problem: people often decline to label obviously moral judgments as moral. This paper explores two possible theories about why people make these "incorrect" judgments: (A) people who make these judgments have less liberal arts education, and so they aren't familiar with these conceptual distinctions, and (B) these unorthodox classifications are caused by a less reflective cognitive style that prevents some people from learning or accessing the distinctive conceptual features of moral judgment. Surprisingly, (A) was disconfirmed: people with

more reported liberal arts education were less likely to classify issues the way that researchers expect. However, (B) was supported: people with higher levels of cognitive reflection most matched researchers' intuitions. Though it runs counter to the self-conception of these disciplines, more reflective cognitive style is also negatively associated with liberal education.

Epistemic mental states elucidated through electrophysiological responses. An experimental approach to Philosophy of Mind

Ricardo Augusto Perera, Monash University, ricardo.perera@monash.edu

Sofia Inês Albornoz Stein, Experimental Philosophy and Cognitive Science Laboratory at Unisinos, siastein@me.com

Adriano Naves De Brito, Experimental Philosophy and Cognitive Science Laboratory at Unisinos, andebrito@me.com

Handout: <https://we.tl/t-p3Nqr0kfJ7>

Keywords: Functionalism, Mental States, Semantics, Electrophysiological Responses, ERP Technique, N400

In the present experiment, we used electrodes to capture electrophysiological signals to detect the existence/absence of beliefs by measuring how they modulate responses using the ERP (Event-Related Potentials) technique and dispensing with the verification of behavioural outputs. Additionally, we obtained empirical data about how these mental states integrate more complex cognitive processes when one reads sentences related to these same mental states. Using the ERP technique, which uses EEG data, we expected to find patterns of electrophysiological responses that “carry information” about the possession or absence of epistemic mental states while participants were reading philosophical sentences ending truly or falsely. Electrodes of primary interest were those located in the central and parietal regions. The ERP component of interest was the N400 (negative deflection whose latency occurs about 400 milliseconds after a stimulus), which is mainly associated with the difficulty of integrating the meaning of a word into the context in which it appears. From the mere brain response satisfying semantic expectations (attenuated N400) or violating them (strong N400), it is possible, we argue, to understand the causal role of the underlying epistemic beliefs modulating the signals and how they are recruited during reading to predict the following lexical items.

Epistemic Side-Effect Effect: A Meta-Analysis

Bartosz Maćkiewicz

Katarzyna Kuś

The epistemic side-effect effect (ESEE; Beebe, Buckwalter 2010) reveals that moral valence of a proposition plays a prominent role in attributions of knowledge. This finding appears to present a challenge to the received epistemological theories.

Many follow-up studies investigated the scope and magnitude of the original findings. ESEE turned out to be robust and not limited to English. It holds for multiple scenarios other than Knobe's original stories and extends to other epistemic concepts. It was also argued that ESEE is independent of the method of gathering responses.

The main aim of our presentation is to assess the existence and magnitude of ESEE. We identified several common types of changes to the experimental design in the literature. Using meta-regression, we analyzed the impact of the study-level covariates on ESEE effect size. Our initial analysis indicates that overall effect size for ESEE is smaller than its counterpart related to intentionality attribution (i.e. Knobe effect). We also found that several factors seem to have a negative impact on the magnitude of the effect size (non-dichotomous scale, gettierization, non-English language, manipulation of the probability of the side effect) while also explaining a significant amount of between-studies variation.

References

- Beebe, J. R. (2013). A Knobe Effect for Belief Ascriptions. *Review of Philosophy and Psychology*, 4, 235- 258.
- Beebe, J. R. (2016). Do Bad People Know More? Interactions between Attributions of Knowledge and Blame. *Synthese*, 193, 2633-2657.
- Beebe, J. R., & Buckwalter, W. (2010). The Epistemic Side-Effect Effect. *Mind and Language*, 25, 474- 498.
- Beebe, J. R., & Jensen, M. (2012). Surprising connections between knowledge and action: The robustness of the epistemic side-effect effect. *Philosophical Psychology*, 25(5), 689-715.
- Beebe, J. R., & Shea, J. (2013). Gettierized Knobe Effects. *Episteme*, 10(3), 219-240.
- Buckwalter, W. (2014). Gettier Made ESEE. *Philosophical Psychology*, 27(3), 368-383.
- Dalbauer, N., & Hergovich, A. (2013). Is What is Worse more Likely? The Probabilistic Explanation of the Epistemic Side-Effect Effect. *Review of Philosophy and Psychology*, 4, 639-657.
- van Dongen, N., Colombo, M., Romero, F., & Sprenger, J. (2020). Intuitions about the reference of proper names: A meta-analysis. *Review of Philosophy and Psychology*, 1-30.
- Borenstein, M., Hedges, L. V., Higgins, J. P., & Rothstein, H. R. (2011). *Introduction to meta-analysis*. John Wiley & Sons.
- Lipsey, M. W., & Wilson, D. B. (2001). *Practical meta-analysis*. SAGE publications, Inc.
- Paprzycka-Hausman, K. (2020). Knowledge of consequences: an explanation of the epistemic side-effect effect. *Synthese*, 1-34.
- Turri, J. (2014). The Problem of ESEE Knowledge. *Ergo*, 1(4), 101-127.

Examining Evaluativity in Legal Discourse: A Comparative Corpus-Linguistic Study of Thick Concepts

Pascale Willemsen, University of Zurich, Department of Philosophy, pascale.willemsen@philos.uzh.ch, <https://www.pascalewillemsen.com/>

Lucien Georges Baumgartner, University of Zurich, Department of Philosophy, lucien.baumgartner@philos.uzh.ch, <https://lucienbaumgartner.github.io/>

Severin Frohofer, University of Zurich, Department of Philosophy, severin.frohofer@uzh.ch

Kevin Reuter, University of Zurich, Department of Philosophy, kevin.reuter@uzh.ch, <http://www.kevinreuter.com/>

How evaluative are legal texts? Do legal professionals speak a more descriptive or perhaps a more evaluative language? In this paper, we present the results of a corpus study in which we examined the use of evaluative language in both the legal domain as well as public discourse. For this purpose, we created two corpora. Our *legal professional corpus* is based on court opinions from the American Court of Appeals. We compared this professional corpus to a *public corpus*, which is based on blog discussions on the internet forum Reddit. While many linguistic phenomena can give insights into evaluativity, we investigated the use of a wide selection of evaluative adjectives (more specifically, thick adjectives) to gain a more comprehensive picture of the degree of evaluativity in the legal domain. Our analysis shows that legal professionals use thick terms less evaluatively than laypeople, which suggests that legal texts are less evaluative than ordinary discussions.

Fragmented and Conflicted? Folk Beliefs about Vision

Eugen Fischer, University of East Anglia, e.fischer@uea.ac.uk, <https://eugenfischer.net/>

Keith Allen, University of York, keith.allen@york.ac.uk

Paul E. Engelhardt, University of East Anglia, p.engelhardt@uea.ac.uk

Belief fragmentation is perhaps one of the philosophically most important cognitive phenomena philosophers have ignored or even outright denied exist: Different cognitive processes, operating under different conditions, generate conflicting beliefs, which are never systematically screened for coherence and stored at different locations in long-term memory, in different ‘belief fragments’ that are subsequently activated by different stimuli.

This paper suggests that conflicts between fragmented beliefs are at the root of philosophical problems including the influential ‘problem of perception’. Approaches to this problem take for granted that ‘Direct Realism’ is the common-sense view of perception. But belief fragmentation challenges the presupposition that there is any such thing as ‘*the*’ common-sense view of perception.

We present a survey that examines whether folk beliefs about vision are fragmented. 100 participants rated their agreement with concrete examples and images representing core claims associated with either Direct Realism or competing Indirect Realism. Exploratory factor analysis identified the content of the corresponding naïve theories. Purely data-driven (K-means) cluster analyses followed by manual cluster analyses identified proportions of participants holding one or both of these incompatible theories. Findings provide evidence of belief fragmentation and suggest that only a small minority hold coherent beliefs about vision.

Gratuitous Risk: perceived danger and recklessness judgements about outdoor sports participants

Philip Ebert, Division of Law and Philosophy, University of Stirling, Scotland, UK, Scientific Member of Behavioural Science Group, School of Management, University of Stirling, Scotland, UK, p.a.ebert@stir.ac.uk, www.philipebert.info, www.varietiesofrisk.info

Ian Durbach, Centre for Research into Ecological and Environmental Modelling, University of St Andrews, Scotland, UK and Centre for Statistics in Ecology, the Environment, and Conservation, University of Cape Town, South Africa, indurbach@gmail.com

Preprint: <https://www.psycharchives.org/handle/20.500.12034/3951>

In this paper, we investigate perceived danger and recklessness judgements (N = 2060) about risk-taking in different outdoor sports ranging from mundane activities to more adventurous activities and discuss our findings within the wider philosophical debate about the notion of transformative experience which could help to explain variations in risk judgements, as well as recent philosophical discussion about non-probabilistic notions of risk and danger, such as the modal and the normic account of risk & danger (Williamson, 2007; Pritchard, 2016; and Ebert, Smith, & Durbach, 2020). The main findings of our study are that substantial differences in fatality rate or in factors with direct links to fatality rates, such as expertise, have far less of an effect on danger, recklessness, and insurance judgements than the kind of sport that is being pursued. Other contextual factors primarily influence recklessness rather than risk judgments, and the most important of these relate to whether there are external parties such as dependants or charity organisations that might come to harm or benefit as a result of the sports participant's activity. These results have philosophical consequences: significant differences in danger and recklessness judgements, despite identical fatality rate, provide some initial support for non-probabilistic conceptions of risk and can thus inform a thriving philosophical debate about different conceptions of risk. Moreover, divergences between danger and recklessness judgement in our survey suggest that the latter is a more complex context-sensitive notion with a distinct moral dimension.

Guilty Artificial Minds

Michael Stuart, Centre for Philosophy of Natural and Social Science, London School of Economics and Social Science; Carl Friedrich von Weizsäcker Centre, University of Tübingen; Centre for Philosophy of Science, University of Geneva, mike.stuart.post@gmail.com, <http://www.michaelstuart.com>

Blameworthiness requires mental states such as intentions, desires and beliefs, which AI systems do not possess. And yet, people blame AI systems anyway. Why? Perhaps people use different notions of blame for different types of agent. Specifically, perhaps people assign blame to humans depending on what that agent knew and intended, and to non-human agents based merely on whether they are the physical cause of some harm. To investigate, we created a scenario in which an agent takes a risk that could lead to a harm. We varied the type of agent in the scenario between a human, a corporation, and an AI system (N = 513), and then asked whether the agent deserved blame, and whether the agent “knew” what it was doing, and desired that outcome. We expected the non-human agents to be blamed despite not being attributed inculpatory mental states. However, we found no differences in attributions of blame across agent type, and, surprisingly, no differences in attributions of inculpatory mental states either. Further, in the case of the AI system, when given the opportunity to downgrade their attribution of knowledge to a more metaphorical or less “mentalistic” phrasing, most participants

did not. It seems, then, that the same sense of blame is being applied across all three kinds of agent. There was one important difference, however, between AI systems on the one hand, and humans and corporations, on the other. AI systems are blamed less when they cause more harm. One possible explanation for this is that when very serious harm is caused, participants try to locate a human who is “really” to blame.

How Do People Balance Death against Lesser Burdens?

Veronika Luptakova, Department of Psychological and Behavioural Science, The London School of Economics and Political Science (LSE), v.luptakova@lse.ac.uk

Alex Voorhoeve, Department of Philosophy, Logic and Scientific Method, LSE, a.e.voorhoeve@lse.ac.uk

Matteo Galizzi, Department of Psychological and Behavioural Science, LSE, m.m.galizzi@lse.ac.uk

Our research provides empirical evidence on moral judgments of the public in the context of distributive justice in healthcare.

First, the paper quantifies what share of the public holds such views on the allocation of health care resources that are inconsistent with standard principles of distributive justice. Second, it experimentally investigates whether *status quo bias* is present in these moral judgments.

We conducted an online survey with an embedded experiment, using a representative sample of the UK population (n = 389). Subjects responded to priority-setting dilemmas about the allocation of healthcare resources, using the person trade-off method. Our treatments consisted in experimentally varying the *status quo* (nudging towards the *utilitarian* or *maximin* principles).

Contrary to standard priority setting principles, around half of our subjects do not always adhere to prioritization by severity. Among those that do, many subjects prioritize saving a small number of people from a severe ailment over saving a multitude from a minor harm, no matter how large this multitude, again contrary to standard principles. We find no statistically significant main treatment effect of variations in the status quo. However, we find statistically significant interaction between treatment and gender, with male subjects being much more susceptible to status quo bias than female subjects. We discuss potential explanations for these results and their implications for policymakers.

Law, Coercion, and Folk Intuitions

Lucas Miotto, University of Maastricht

Guilherme de Almeida, PUC-Rio

Noel Struchiner, PUC-Rio

Long abstract:

<https://drive.google.com/file/d/1pdqsK9NqqUq1nCMdXEHWd5S6rPX68B3g/view?usp=sharing>

Keywords: Experimental Jurisprudence, Legal Philosophy, Analytical Jurisprudence, The Nature of Law, Coercion

Up until recently, the almost consensual view in analytical jurisprudence about the relationship between law and coercion has been that legal systems are only contingently coercive. Recently, however, this traditional view has been challenged. The challenge is motivated by doubts about the reliability of the thought experiments that support this view: thought experiments involving sanction-less legal systems in a society of angels. Critics have not only argued that intuitions about such thought experiments do not bear on our understanding of *actual* legal systems, but also that “the man on the Clapham Omnibus” – as Frederick Schauer puts it – would exhibit significantly different intuitions about such scenarios. Laypeople would not hold that there is law in a society of angels because the view that law is *necessarily* coercive “enjoys widespread support among laypersons” – or so critics claim (Kenneth Himma, *Coercion and the Nature of Law* (Oxford University Press 2020) viii.). This is obviously an empirical assumption about laypeople’s intuitions. Critics, however, never systematically polled the “man on the Clapham Omnibus”. We boarded that bus.

We conducted three pre-registered studies to investigate laypeople’s intuitions about the relationship between law and coercion. Contrary to what critics believe, our results by and large support the traditional view in analytical legal philosophy.

Morality Rules: Rule Application is Shaped by Moral Evaluation

Brian Flanagan, Brian.Flanagan@mu.ie

Ivar Hannikainen, ivar.hannikainen@gmail.com

Noel Struchiner, struchiner@gmail.com

Guilherme de Almeida, almeida2808@gmail.com

Formal rules organize large parts of modern social life, from the doings of the individual household to the conduct of the inhabitants of the nation state. But there has been a lack of empirical investigation into the nature of the rule itself, and theoretical consensus remains elusive. This paper applies the experimental method to the issue of whether rules intuitively consist in the historical features of their formulation alone, notably, in their text or historical purpose, or in some combination of those features and moral values. We report 6 studies (N 2,013) that indicate that the perceived morality of an individual’s conduct and beliefs is a major predictor of judgments of whether she has violated a rule, especially of judgments made heuristically, under time pressure. Equally, we found that morality’s influence on rule application increases when the relevant agent is ascribed knowledge of the rule’s purpose. Finally, in line with this result, we found that, contrary to the conventional theoretical understanding, an agent’s ignorance of a rule’s existence is an intuitively significant exculpatory factor in rule application.

Outcome Effects, Moral Luck and the Hindsight Bias

Markus Kneer, Guilty Minds Lab Zurich, markus.kneer@uzh.ch, <https://www.markuskneer.com>,
<https://www.guiltymindslab.com>

Izabela Skoczeń, Guilty Minds Lab Zurich, izaskoczen@gmail.com,
<https://izabelaskoczen.wordpress.com>, <https://www.guiltymindslab.com>

Preprint: <https://bit.ly/3vJeNDR>

In a series of ten preregistered experiments (N=2043), we investigate the effect of outcome valence on judgments of probability, negligence, and culpability – a phenomenon sometimes labelled moral (and legal) luck. We found that harmful outcomes, when contrasted with neutral outcomes, lead to increased perceived probability of harm ex post, and consequently to increased attribution of negligence and culpability. Rather than simply postulating a hindsight bias (as is common), we employ a variety of empirical means to demonstrate that the outcome-driven asymmetry across perceived probabilities constitutes a systematic cognitive distortion. We then explore three distinct strategies to alleviate the hindsight bias and its downstream effects on mens rea and culpability ascriptions. Not all are successful, but at least some prove promising. They should, we argue, be taken into consideration in criminal jurisprudence, where distortions due to the hindsight bias are likely considerable and deeply disconcerting.

Outcome-valence and the reasonable person standard

Markus Kneer

Philosophical Expertise Put to the Test

Samuel Schindler, Centre for Science Studies, Aarhus University, samuel.schindler@css.au.dk,
www.samuelschindler.org

Pierre Saint-Germier, Institut supérieur de philosophie, University of Louvain, Belgium,
saintgermier.p@gmail.com, <https://uclouvain.academia.edu/PierreSaintGermier>

The so-called expertise defence against sceptical challenges from experimental philosophy has recently come under attack: there are several studies claiming to have found *direct* evidence that philosophers' judgments in thought experiments, too, are susceptible to erroneous effects. In this paper we distinguish between the customary "immune experts" version of the expertise defence and an "informed experts" version. On the *informed expertise defence*, we argue, philosophers' judgments in thought experiments could be preferable to the folks', *even if* it were true that philosophers' judgments are no less immune to confounders than the folk's. We present results from an experimental study comparing philosophers and non-philosophers (m=484), which, we argue, supports the "informed experts" version of the expertise defence.

Psychological Implications of Gendered AIs and Their Social Impact

Javier Gomez-Lavin, The University of Pennsylvania, jgomezlavin@gmail.com,
<https://jgomezlavin.com>

Hannah Read, Duke University, hannahread01@gmail.com, <https://hannah-read.com>

Andrea Beltrama, The University of Pennsylvania, andrea.beltrama@gmail.com,
<https://andreabeltrama.com/>

Lisa Miracchi, The University of Pennsylvania, miracchi@gmail.com, <http://www.lisamiracchi.com/>

Widespread reliance on feminized AIs runs the risk of perpetuating unjust gender norms regarding women's subservience. Yet, despite the potential impacts of feminized AI on gender norms and oppression, these issues remain largely unaddressed. We therefore conducted two studies to assess how expectations differ depending on the stereotypical gendering of the AI assistant and task.

Study 1 normalized and determined the gender typicality of 52 common-place tasks associated with 13 activities involving gendered AIs. Each participant was first presented with 26 activity-task combinations. Participants then rated each task's gender typicality, complexity, intellectual and emotional demands, as well as their familiarity with the task. Tasks differed significantly across gender typicality.

Study 2—a between-subjects 2x2 design—randomly assigned participants to one of four conditions featuring vignettes involving an AI with a **feminine** or **masculine name** and **pronouns** to help perform a **feminine** or **masculine task**, as determined by our first study. Participants were then instructed to rate the acceptability of the AI assignment and the AI's competence, agency, and emotional warmth.

Results from both studies provide a foundation for further empirical work and motivate the need for careful empirical study of these issues.

'Really Know': A Linguistic and Philosophical Analysis

Mark Satta, Wayne State University, mark.satta@wayne.edu, www.marksatta.com

Lacey J. Davidson, California Lutheran University, ljdavidson@callutheran.edu,
www.laceyjdavidson.com

Augustus R. Wachbrit, Stanford University, wachbrit@stanford.edu

Epistemologists have offered different arm-chair views on the following question: How often are we willing to claim that we 'really know'? We use corpus linguistics to test the following two hypotheses. First, competent speakers of English are quicker to ascribe that someone 'knows' than that someone 'really knows'. Second, competent speakers of English are quicker to deny that someone 'really knows' than that someone 'knows'. To test these hypotheses, we ran two studies. In both studies, we took two random samples from a linguistic corpus. The first sample contained sentences with the lexeme 'know' unmodified by 'really'. The second sample contained sentences with the lexeme 'know' modified by 'really'. We independently hand-coded the sentences into categories, including knowledge ascribing, knowledge positive (unnegated use of 'know'), knowledge denying, and knowledge negative (negated use of 'know'). Both studies showed substantially higher levels of knowledge ascribing and knowledge positive sentences in the 'know' condition when compared to the 'really know' condition. Both studies also showed substantially higher levels of knowledge denying and

knowledge negative sentences in the ‘really know’ condition when compared to the ‘know’ condition. Thus, both studies support our hypotheses.

Revisionist attitude reports and inference patterns

Tomasz Zyglewicz, The Graduate Center, CUNY, tzyglewicz@gradcenter.cuny.edu

Keywords: x-phi, attitude reports, belief reports, desire reports, inference patterns

In a recent paper, Kyle Blumberg and Harvey Lederman put forward a revisionist semantics for attitude reports, according to which the truth of attitude ascriptions is systematically dependent on the information unavailable to the subject (Blumberg and Lederman 2021). They motivate the view with scenarios like the following:

Tennis. Ann is a six-year-old girl whom Pete, an expert in tennis pedagogy has never met and whose existence he is unaware of. Pete believes that every six-year-old can learn to play tennis in ten lessons. Jane says to Ann’s father, Jim:

(1) Pete believes Ann can learn to play tennis in ten lessons.

According to their theory, (1) is true because it is a result of revising Pete’s belief state in the light of information available to Jim and Jane.

B&L take this phenomenon of revisionist reporting to be fairly general: it extends across various attitude verbs and inference patterns. However, I ran a quick explanatory study that revealed robust differences in the folk’s intuitions about two cases invoked by B&L. In this talk I will present the results of a more rigorous follow-up study. A 2x2 design will be used, where the two variables are the attitude verb (‘believes’ vs ‘wants’) and the inference pattern (universal instantiation vs disjunction elimination). I hypothesize that the inference pattern influences the acceptability of revised attitude reports more than the attitude verb does.

References: Blumberg, Kyle, and Harvey Lederman. 2021. “Revisionist Reporting.” *Philosophical Studies*, no. 178: 755–83. <https://doi.org/10.1007/s11098-020-01457-4>.

Should I say that? An experimental investigation of the norm of assertion

Neri Marsili, (1) Alma Mater Studiorum – Università di Bologna, (2) Universitat de Barcelona, nerimarsili@gmail.com, <https://nerimarsili.weebly.com>

Alex Wiegmann, Ruhr-Universität Bochum, Alex.Wiegmann@psych.uni-goettingen.de, <https://www.psych.uni-goettingen.de/de/cognition/team/wiegmann>

Preprint: <https://psyarxiv.com/cs45j/>

Assertions are our standard communicative devices for sharing and acquiring information. Recent studies seemingly provide converging evidence that assertions are subject to a *factive* norm: you are entitled to make an assertion only if it is true. However, these studies assume that we can treat participants’ judgements about what an agent ‘should say’ as evidence of their intuitions about assertability. We argue that this assumption is incorrect, so the conclusions drawn in the aforementioned studies are unwarranted. We provide evidence that most people do not interpret statements about what one ‘should say’ as statements about assertability, but rather as statements

about what is in the agent's interest to do. Measures for prompting the intended reading of the test question are identified, and their efficacy is tested. We found that when these measures are implemented, people's judgements consistently and overwhelmingly align with non-factive accounts of assertion.

Socratic Questionnaires

Nathaniel Hansen, University of Reading, n.d.hansen@reading.ac.uk,
<https://sites.google.com/site/nathanielhansen/>

Kathryn Francis, University of Keele

Hamish Greening, University of Reading

Paper: <https://docs.google.com/document/d/1k5UT9mPHZyGyoJRI-m-L49NmQkJKTM0VL3Y27WjY9vO0/edit>

When experimental participants are given the chance to reflect and revise their initial judgments in a conversational context, their responses to some philosophical scenarios differ from responses to those same scenarios presented in a traditional static survey. This finding complicates recent findings that various manipulations of reflectiveness have no effect on participants' judgments about philosophical scenarios. Conversational experiments also illuminate debates about how exactly participants are understanding the scenarios they read in experimental philosophy studies, in particular whether they are replacing difficult questions with questions that are more easily answered, and to what extent participants are imagining the scenarios they read in ways that differ from what is explicitly stated by experimenters. We argue that conversational experiments—"Socratic questionnaires"—should be used alongside traditional survey designs in the experimental investigation of philosophical questions.

The Asian Disease Problem and Doing/Allowing Harm

Tomasz Żuradzki, Jagiellonian University in Kraków

Marta Maj, Jagiellonian University in Kraków

Paulina Szwed, Jagiellonian University in Kraków

Keywords: Asian Disease Problem, doing and allowing harm, risk attitudes, framing effect, non-consequentialism

The Asian Disease Problem (ADP), one of the most famous experimental paradigms proposed by Tversky and Kahneman (1981), brings some conceptual similarities to philosophical discussions around the distinction between doing and allowing harm (Woollard 2015). Our research questions are twofold. First, we challenge both – the standard and some non-standard (i.e. highlighting that the options in the original version of ADP are not fully equivalent) interpretations of the ADP. Second, we challenge the views claiming that the allegedly inconsistent choices in ADP may be fully justified by non-consequentialist values, that is, justified risk-seeking when it comes to negative framing and justified risk-aversion when it comes to positive framing (Dreisbach & Guevara, 2017). To challenge the above views, we recruited, in three waves, 1106 participants who completed (just before the recent pandemic) online questionnaires about ADP, and some of them were asked to provide the reasons justifying their decisions to get insight into their motivations. The results show that the ADP does not disappear when we simultaneously presented all four options to participants, what has been implied

by both standard and some non-standard interpretations of ADP. Next, our results have only partially confirmed the non-consequentialist justification of the allegedly inconsistent choices in the ADP.

The Future of Virtual Reality in Experimental Research: Revolutionary or Problematic?

AI and future technologies symposium, invited talk

Kathryn Francis, School of Psychology, Keele University

In recent years, immersive Virtual Reality (VR) has allowed the development of innovative tools and measures for experimental research in philosophy and psychology. These technologies have been invaluable in experimental research by improving ecological validity, generating sensory feedback, and allowing real-time performance recording. Using moral psychology as a case study, I discuss the impact that VR has had on methodological approaches and how it has raised important theoretical questions in the field. I also consider technical and methodological issues concerning the application of VR and related technologies across several research areas. Importantly, and when considering the future of these technologies in experimental research, I highlight the ethical challenges researchers face in using virtual and augmented environments to further their research.

Theoretical virtues: do scientists think what philosophers think they ought to think?

Samuel Schindler, Centre for Science Studies, Aarhus University, samuel.schindler@css.au.dk, www.samuelschindler.org

Theoretical virtues play an important role in several debates in the philosophy of science. This paper presents the results of a quantitative study of the views natural and social scientists hold about the role and nature of theoretical virtues. The results were compared to views held by philosophers. Some of the main results are: (i) there is broad agreement across all three groups about how the virtues are to be ranked in terms of importance, (ii) all groups agree that unification is an epistemic virtue and there is even some evidence that simplicity is viewed as epistemic by scientists, (iii) all groups consider syntactic parsimony as more important than ontological parsimony, (iv) all groups consider unifying power as independent from simplicity, (v) all groups find simplicity and unification aesthetically pleasing, and (vi) scientists regard the aesthetic appeal of unifying power, and in particular, the aesthetic appeal of simplicity as truth-tracking.

Tracing Thick Concepts Through Corpora

Lucien Baumgartner, University of Zurich, Department of Philosophy,
lucien.baumgartner@philos.uzh.ch, <https://lucienbaumgartner.github.io/>

Kevin Reuter, University of Zurich, Department of Philosophy, kevin.reuter@uzh.ch,
<http://www.kevinreuter.com/>

Pascale Willemsen, University of Zurich, Department of Philosophy,
<https://www.pascalewillemsen.com/>

Philosophers and linguists currently lack the means to reliably identify evaluative concepts and to measure these concepts' evaluative intensity. Using a corpus-based approach, we present two new methods that allow us to distinguish evaluatively thick adjectives like 'courageous' from descriptive adjectives like 'tall', and from value-associated adjectives like 'sunny'. In Study 1, we show that evaluative adjectives behave differently from descriptive adjectives when in combination with other adjectives through use of the connective 'and'. Study 2 reveals that the modifiers 'truly' and 'really' highlight the evaluative dimension of thick and thin adjectives, allowing for them to be uniquely classified. Based on the results of our studies, we believe the operationalizations we suggest may pave the way for a more quantitative approach to the study of thick concepts.

Weakness of Will: Existence and Rationality

Julian Fink, University of Bayreuth, julian.fink@uni-bayreuth.de

Kevin Reuter, University of Zurich

Michael Messerli, University of Zurich

Weakness of will remains a perplexing issue. Though philosophers have made substantial progress in homing in on what counts as weak-willed intention and action, there has been little effort to investigate the empirical foundations of weakness of will. In particular, no conclusive evidence for the existence of weakness of will has yet been produced. This paper seeks to provide such empirical support. The second goal of this paper is to study the ir/rationality people often associate with weakness of will. We show that rationality assessments have a dual sensitivity: on the one hand, we can improve our rationality by intending something we believe we ought to do. On the other hand, we can improve our rationality by making our intention and what we value coherent. In exploring whether this dual sensitivity can be accommodated within a unified theory of rationality, we argue that only a success-based account of rationality can do so. Coherence- and reasons-based accounts fail to make sense of rationality assessments.

Western Skeptic vs Indian Realist. Cross-Cultural Differences in the Zebra Case intuitions

Krzysztof Sękowski, Faculty of Philosophy, University of Warsaw, k.sekowski@uw.edu.pl

Adrian Ziółkowski, Faculty of Philosophy, University of Warsaw, adrian.a.ziolkowski@uw.edu.pl

Maciej Tarnowski, Faculty of Philosophy, University of Warsaw, m.tarnowski3@student.uw.edu.pl

In their famous paper, Weinberg, Nichols and Stich (2001) reported cross-cultural differences in epistemic intuitions. However, most of their findings failed to replicate. The exception here is the purported differences between Westerners and Indians in knowledge ascriptions concerning the Zebra Case, which was never properly replicated. Our study replicates the above mentioned experiment. The analysis found a significant difference between the ethnic groups in question: Indians were more likely to attribute knowledge in the Zebra Case than Westerners. In our talk, we offer an explanation of our result that takes into account the fact that replications of Weinberg et al.'s other experiments have failed. We argue that the Zebra Case is unique among the vignettes tested by Weinberg et al. and it should not be regarded as a Gettier case but rather as a scenario exhibiting skeptical pressure concerning the reliability of sense-perception. We argue that skepticism towards perception as a means of gaining knowledge is a trope that is deeply rooted in Western epistemology. However, as the examination of the skeptic scenarios discussed by philosophers of the Indian *Nyaya* tradition and their adversaries shows, it is very much absent from Classical Indian philosophical inquiry.

What are the components of a meaningful life?

Joffrey Fuhrer, University of Geneva (Unige), PhD in philosophy, affiliated to the Swiss Center for Affective Sciences, Joffrey.Fuhrer@Unige.ch, <https://www.unige.ch/cisa/center/members/fuhrer-joffrey/>

My talk is about my study named “several shades of lives” which aims at testing laypeople’s intuitions about what makes a life happy and meaningful, as well as the most important factors that influence these judgments.

In this study, participants were shown five vignettes telling the life (“shades”) of a character and were then asked to answer several questions about how they attribute meaningfulness, happiness, and self-perception of meaningfulness to a life. Each type of vignette differs on some important features, such as the degree of engagement in some activity, the degree of fulfillment in life, the degree of importance on one’s main activity, and one’s degree of skillfulness for that activity.

Our results showed that some lives exhibited very specific patterns in terms of happiness/meaningfulness/self-perception of meaningfulness attributions (H1). We also found that some variables were more important than others for each of these attributions (H2). In an effort to check that attribution of meaningfulness and happiness would be different, we demonstrated (H3) that there was no significant difference between the baseline correlation we hypothesized and the one we found.

What is the folk concept of life?

Claus Beisbart, University of Bern, Department of Philosophy, Claus.Beisbart@philo.unibe.ch, https://www.philosophie.unibe.ch/ueber_uns/personen/beisbart/index_ger.html

Kevin Reuter, University of Zurich, Philosophisches Seminar, kevin.reuter@uzh.ch, <http://www.kevinreuter.com/>

The old question “What is life?” has recently drawn intensified attention. If an answer to the question is to move beyond mere stipulation, it has to keep some similarity to the folk concept of life. In this talk, we present results from experiments on the folk concept of life. Using a semantic feature task (Study 1), we show that people associate macroscopic functionings such as growth, breath or nutrition with living beings. If asked for universal features of life (Study 2), however, participants most often name cells and organic matter. We use the contrast between the salient and the universal features of life to argue that life is a natural kind concept.

When is reappropriation appropriate? Lessons from 'gossip'

Tomasz Zyglewicz, The Graduate Center, CUNY, tzyglewicz@gradcenter.cuny.edu, <https://zyglewicz.com>

Shannon Brick, The Graduate Center, CUNY, sbrick@gradcenter.cuny.edu, <https://www.shannon-brick.com/>

Manuscript: https://www.dropbox.com/s/pepeynv3h5kav50/Zyglewicz%26Brick%20-%20gossip_appropriation_manuscript.pdf?dl=0

Keywords: gossip, reappropriation, reclamation, conceptual engineering, experimental philosophy

It has been recently suggested that gossip is more than just idle talk. In particular, authors working within feminist epistemology have argued that gossip is a relatively safe way for members of marginalized groups to resist the norms and values of their oppressors, as well as circulate potentially protective information (PPI). In this paper we oppose the idea of calling such subversive exchanges “gossip.” Our case consists of a descriptive claim and a normative claim. The descriptive claim is that the folk are not inclined to classify an utterance about an absentee as gossip when it conveys PPI. We substantiate this claim by presenting results of an experiment devised specifically to test it. The normative claim, in turn, is that utterances conveying PPI should not be labeled as “gossip”. We use our argument for the normative claim to mount a more general account of when it is and is not appropriate to reappropriate stigmatizing terms. We call this account the World of Our Dreams Criterion. According to it, a term T is apt for reappropriation if there is no use for a pejorative use of T in the world envisaged by the group targeted by T .