

Conference on Experimental Philosophy

19th September - 21st September 2019
University of Bern

Organizer: Kevin Reuter

**Co-Organizers: Experimental Philosophy Group Germany
Albert Newen, Florian Cova, Markus Kneer, & Pascale Willemsen**

Website: <https://sites.google.com/view/xphibern2019/>

Keynote Speakers:

Chiara Lisciandra (Groningen)
David Rose (Florida State)
Eric Schwitzgebel (California, Riverside)
Christina Starbans (Toronto)
Brent Strickland (Institute Jean Nicod)
Jennifer Cole Wright (Charleston)



Design: Pascale Willemsen

FNSNF

FONDS NATIONAL SUISSE
SCHWEIZERISCHER NATIONALFONDS
FONDO NAZIONALE SVIZZERO
SWISS NATIONAL SCIENCE FOUNDATION

u^b

UNIVERSITÄT
BERN

ABSTRACTS

Abstracts of Keynote Talks.....	p. 3-8
Abstracts of Contributed Talks.....	p. 9-33
Abstracts of Poster Presentations & Blitz Talks.....	p. 34-44

(Norms)²: Norms about Norms

Chiara Lisciandra

In this talk, I outline and defend the view that variations in compliance levels with one and the same norm represent different norms about following norms. In support of this claim, I will first argue that classic game-theoretic accounts, which define norms as Nash equilibria of noncooperative games, typically consider variations in compliance levels as separate norms. After that, I suggest a more fine-grained, game-theoretic distinction that accounts for degrees of compliance with the same norm and show how to incorporate such an account in a psychological framework. Finally, the paper examines what given degrees of compliance can reveal about the dynamics underlying the process of norm change. I will argue that they are indicators of different reactions to the introduction of new norms.

The Depth of Teleology

David Rose

For roughly the past thirty years, the dominant view of psychological essentialism is that essence is represented by a placeholder. This view—advocated by Susan Gelman, Frank Keil, Douglas Medin and Andrew Ortony—is Lockean: an essence is the underlying property that causes observable features, though we may not know what that underlying property is. But a range of research indicates that teleology impacts a broad range of judgments, including judgments of whether some parts compose a whole object (Rose & Schaffer, 2017) and whether an object persists through part alterations (Rose, 2015; Rose, Schaffer & Tobia, 2019). This might be due to a deeper fact about teleological thinking; namely, that we essentialize categories in terms of teleology. Using the same procedures that are typically used to provide evidence of essentialist thinking, I'll provide evidence that people essentialize a broad range of categories in terms of teleology. I will then trace out the implications for philosophical and psychological theorizing on categorization and concepts arguing, for instance, that teleological or Aristotelian essentialism challenges a recent Lockean/Platonic view of essence which is inspired by work on dual-character concepts.

Do Philosophical Arguments Influence Moral Behavior? Data on Meat Ethics and Charitable Giving

Eric Schwitzgebel

Do philosophical arguments influence real-world moral behavior? In this talk:

- (1.) I review my empirical findings on the not-especially-ethical behavior of ethics professors.
- (2.) I present new data from a study in which students read a philosophy article defending vegetarianism and discussed the material in groups for 50 minutes. Compared to a control group, students who read and discussed the article were more likely to say eating factory-farmed meat was unethical. We then obtained meal card purchase receipts from campus dining services, which confirmed that students actually purchased less meat after the reading and discussion.
- (3.) I present new data from two studies in which MTurk workers were exposed to various arguments for charitable giving, then given a surprise bonus and an option to donate some portion of that bonus to charity. Donation rates were no higher than in the control condition. However, workers who read an emotionally moving narrative did donate at higher rates than in the control condition.

How Temptation Makes Us Moral

Christina Starbans

tba

Re-thinking "core cognition": Core physics in adult vision and its theoretical implications

Brent Strickland

“Core physics” refers to representations of the physical world that appear early in cognitive development, universally across human cultures, and are likely a result of our genetic endowment. In practice the majority of the work on core physics has come from infant psychology, which has compellingly demonstrated over nearly 40 years that very young, pre-verbal infants possess rich representations of objects and types of relationships between objects. The current talk focuses on a less studied aspect of this part of core cognition: it operates automatically and sometimes unconsciously in adult visual perception. I illustrate this with a few recent experimental studies from my lab and others, and I argue that the insights from adult vision cast reasonable doubt on some long standing theoretical commitments from developmental psychology about the place of core cognition in cognitive architecture. In particular we should not think of core cognition as being (primarily) cognitive! Instead many infant results appealing to precocious "reasoning" abilities (about objecthood or causality) are better explained through mechanisms that belong in perception proper and are shared by infants and adults. In addition to this somewhat deflationary perspective there is a positive upshot: perception is actually far more sophisticated than we typically give it credit for.

Humility, Moral Realism, Conviction, and Tolerance – Navigating a Morally Imperfect World

Jen Cole Wright

What is acceptable (even desirable) diversity – and when does that diversity become deviance? This question is critical for creating and maintaining healthy socio-cultural normative structures (i.e., socially ‘normed’ beliefs, values, practices, behaviors, etc.) that allow individuals within a culture to function well (and ideally thrive). The tension generated here introduces critical space for variation, both within and between communities. It also highlights a problem—the imperfection of our moral knowledge and the vulnerability of our normative structures to error and corruption. The challenge we face is having an understanding of ourselves as moral beings pursuing “the good life” that is stable enough to be meaningfully shared and passed down to future generations, yet flexible enough to adapt and change as our shared experiences bring that understanding into question. Moral conviction has a critical psycho-social role to play in this endeavor—and it is a paradoxical role, insofar as it is necessary both for protecting existing normative structures from corruption and for spearheading corrective endeavors, when normative structures have become dysfunctional and change is required. Of course, this is precisely what makes conviction a troublesome bedfellow—given the imperfection of our moral knowledge, it always runs the risk becoming unreasonably dogmatic and oppressive, on the one hand, and irrationally rebellious, on the other. How do we utilize moral conviction to our benefit, while avoiding its dangers? I will consider this question—and in particular, the importance of humility (specifically) and virtue (more generally) in helping us to wield conviction wisely.

Cross-Linguistic Evidence of Essentialist Beliefs about the Law

Ivar R. Hannikainen, Kevin Tobia, Vilius Dranseika, Fernando Aguiar, Guilherme Almeida, Bartosz M. Janik, Maciej Próchnicki, Piotr Bystranowski, Alejandro Rosas, Niek Strohmaier

Laws differ widely from country to country and have evolved dramatically throughout history. Despite such pervasive variation, many legal theorists have argued that there is an essence of law. For instance, legal philosopher Lon Fuller famously highlighted eight properties (e.g., prospectivity, publicity, or intelligibility) that all legal systems share and in virtue of which they are law. Among anthropologists, the concept of law has been also treated as a plausible human universal. In the present work, we were interested in understanding whether people around the world share intuitive beliefs about the nature of law. To answer this question, we drew on a novel paradigm of essentialist thought: Some participants were asked whether laws exhibit a given property (e.g., whether all laws are prospective or some are retrospective), and we referred to these judgments as empirical judgments. Other participants were asked whether laws would “have to” exhibit a given property (e.g., whether all laws have to be prospective or some could be retrospective), which we refer to as necessity judgments. In a between-subjects design, we asked participants to make either empirical judgments or necessity judgments regarding each of eight procedural principles about the law. Participants in seven different countries ($N = 2,179$) believed that (P1) laws “have to” observe a variety of procedural principles even though (P2) laws in practice often do not. We wondered whether statements that a law “could not” violate certain principles were interpreted as modals of obligation instead of necessity: A post-test question revealed that the perceived objective of the task was to describe “the requirements for something to count as law”, and not “how laws should be, according to [participants] beliefs about right and wrong”. Thus, our results indicate that laws are believed to have necessary properties that legal systems violate in practice. We conclude that people share robust intuitive beliefs about the nature of law that (i) defy their ordinary empirical grasp of how legal systems function, and thus (ii) are unlikely to emerge from experience alone.

Remembering Conceptually Entails Knowing

Alexandra Nolte and David Rose

Does remembering conceptually entail knowing? Many philosophers endorse the epistemic theory of memory, arguing that remembering is a kind of knowing (e.g., Adams, 2011; Locke, 1971; Moon, 2013; Audi, 2003; Malcolm, 1963; Margalit, 2002; Anscombe 1981b; Williamson, 2000). More specifically, semantic memory requires knowledge. On this view, remembering that *p* requires everything that is required for knowing that *p*.

One of the main justifications for endorsing the epistemic theory of memory is that it fits our ordinary use of “remembers” and “knows.” Yet, some maintain that there are intuitive counterexamples to the epistemic theory of memory (e.g., Martin and Deutcher, 1966; Bernecker, 2008, 2010; Lehrer and Richard, 1975). These philosophers claim that one can remember that *p* without believing that *p*, without having a justified belief that *p*, or without having a non-accidental true belief that *p*. However, proponents of the epistemic theory of memory dispute the counterexamples, claiming that the folk surely wouldn’t judge that one can remembering that *p* without knowing that *p*.

So we see a dispute over whether the folk think that remembering requires knowing with wide disagreement over what the folk think. This seems like an opportunity to provide empirical input into the debate. We thus ran four studies—all within subject designs—aimed at investigating whether remembering conceptually entails knowing.

We first consider whether people operate with a factive conception of remembering and knowing. While some research indicates that people have a factive conception of knowledge (Buckwalter, 2014), other recent work indicates that people may have a non-factive conception of remembering (e.g., Dranseika, forthcoming). However, we argue that non-factive ascriptions of remembering may be due to perspective taking, where a participant imagines what is true from the protagonist’s point of view. Our results, using three different storylines and varying in whether *p* is true or false, indicate that when controlling for perspective taking, people operate with factive conception of both knowing and remembering. Importantly, they also show that people deny remembering and knowing at similar rates when the target proposition is false. And they ascribe remembering and knowing at similar rates when the target proposition is true. Almost no one ascribes remembering to a protagonist while at the same time denying knowledge.

The received view of knowledge is that knowledge entails belief. So if remembering entails knowing, it follows that remembering entails belief. Yet Martin and Deutcher (1966) present what is perhaps the most famous case of remembering without belief. We tested this case along with two others, including Colin Radford’s famous Unconfident Examinee case. And we also included matched controls where the protagonist clearly believed the target proposition. Controlling for perspective taking, as in our first study, we find that people overall deny both remembering and knowing when denying belief. When they ascribe belief, they overwhelmingly attribute both remembering and knowing.

Next we turn to justification. Bernecker (2008) presents a case where an individual forms a justified true belief that *p* but then is presented with misleading evidence that is not defeated.

He maintains that the presence of the undefeated defeater rules against knowing that p. Nonetheless, the protagonist remembers that p. We tested this case, along with two other stories, and also included control cases where the belief was clearly justified. Controlling for perspective taking, we find that when people judge that a protagonist's belief is unjustified, they deny both remembering and knowing; when they judge that a protagonist's belief is justified, they ascribe both remembering and knowing.

One final type of case remains: one that involves a justified, true belief that is accidentally true. A range of findings suggest that people deny knowledge when a protagonist has been

Gettiered (Machery, Stich, Rose, et al., 2015; Machery, Stich, Rose et al., 2017a; 2017b; Turri, 2013). But perhaps people, as Bernecker (2008; 2010) suggests, deny knowledge when a protagonist has a justified, true belief that is accidentally true, yet claim that the protagonist remembers the relevant proposition. Again, here we found people tended to either ascribe both remembering and knowing or deny both remembering and knowing.

Our findings indicate that, on the ordinary view, remembering requires knowing. They thus build on an independently plausible view about knowledge, namely that it requires non-accidentally true, justified belief and add the surprising addition that remembering requires everything required for knowing. We take our findings to serve as impressive evidence favoring the epistemic theory of memory and conclude by tracing out the consequences for the philosophical debate over the epistemic theory of memory.

Intuitions on the Individuation of Musical Works

Elzė Sigutė Mikalonytė

Philosophers often emphasize the intuitiveness as an advantage of their theories of ontology of musical works. Despite many incompatible predictions of what folk intuitions on relevant questions might be, so far there is only one experimental philosophy study on the repeatability of musical works (Bartel 2017). In order to investigate listeners intuitions on the individuation of works of classical music and to find out which properties of musical performance have to change for listeners to consider that the musical performance ceased to be of the same musical work, we created seven scenarios, which reflect the main disagreements between positions held by musical ontologists, including pure and timbral sonicism, instrumentalism, and contextualism. All scenarios describe two musical performances which are: (a) two identically sounding performances of two identical scores which were independently created by two composers; (b) two identically sounding performances of two identical scores which were created by two different composers using the same technique – rewriting backwards an already existing work; (c) differently sounding musical performances of two different scores written by different composers when one score is a reversed version of another; (d) two performances different only in respect to emotional expressivity; (e) different only in instrument and timbre; (f) different only in instrument, but not timbre; (g) different only in respect to the images evoked in the listeners. Altogether 445 people either with or without music education participated in two studies. In the first study, each participant (N = 249) was asked to evaluate each scenario by choosing one of the answer options: (1) Definitely the same work; (2) Probably the same work; (3) I don't know / I can't tell; (4) Probably two distinct works; (5) Definitely two distinct works. In the second study, participants (N = 196) could choose between these answers: (1) In both cases one and the same musical work is being performed; (2) In first case one musical work is being performed and in the second – another; (3) It can be said both that one and that two distinct musical works are being performed; (4) I can't say whether one work or two works are being performed. Our results across two studies show that in most of the cases there exist prevalent intuitions on questions of the individuation of musical works. In most cases, participants had clear intuitions whether scenarios speak of one or about two musical works. The answer option in the second study that both these answers are correct was less popular, and the agnostic answer was even less popular in all of the scenarios. The results also show that emotional expressivity, instrument, timbre, and images evoked in the listeners were not considered as properties individuating musical works. However, the musical works were held to be different if the composers were different. In most cases, intuitions were of the same nature but more strongly expressed among those who have musical education. The main finding is that pure sonicism complemented with additional significance of the composer's creativity seems to be the most intuitive position.

Reports from Twin Earth

Jussi Haukioja, Mons Nyquist and Jussi Jylkkä

In this paper, we present results from three experiments on ordinary speakers' usage of natural kind terms (NKTs), partly using cases modeled after Putnam's Twin Earth. The Twin Earth case is widely taken to provide strong support for an externalist view that combines a causal-historical theory of reference with essentialism: sharing underlying nature with standard samples of a kind is necessary and sufficient for belonging in the extension of the relevant NKT. Previous experimental work on NKTs has cast some doubt on externalism, but the results are far from conclusive. We presented subjects with two types of scenarios (for 5 NKTs): Twin Earth cases, where new samples share appearance, but not underlying structure, with the standard samples, and reverse Twin Earth cases, where new samples share underlying structure, but not appearance, with standard samples. The subjects' usage was probed both with an elicited production task and with forced-choice questions. In addition, we attempted to see to what extent the subjects were inclined to defer to the relevant experts. In the first two experiments, the underlying nature or the appearance of the new samples was radically different from the standard samples, as in Putnam's Twin Earth case. The standard Putnamean judgments concerning Twin Earth cases (for "water", as well as for other NKTs) were confirmed, but at the same time the subjects were systematically unwilling to apply a NKT to samples which shared an underlying structure, but not their appearance, with the standard samples. For example, watery XYZ was not categorized as water, but neither was solid and smelly H₂O. This suggests that speakers take both sharing of underlying structure and sharing appearance with standard samples of a kind to be necessary for belonging in the extension of the relevant NKT. This result is quite problematic both for mainstream externalist and mainstream internalist theories of NKTs. These results do not, however, tell us whether speakers require new samples to be identical to standard samples in underlying structure and appearance, in order for them to belong to the familiar kind, or whether some degree of similarity in one or both respects would be enough. To look at this issue, we ran a third experiment, using scenarios similar to the above. However, now the underlying natures and appearances were not as radically different as in the first two experiments. Two less radical versions of each scenario were prepared, one where the appearance/underlying nature was only slightly different from the standard samples, and another where these were clearly different, but not as radically different as in the Twin Earth case. The data from the third experiments reveals a gradualness in the subjects' answers: the more the underlying structure or the appearance differs from those found in the standard samples, the less likely subjects are to categorize new samples as belonging to the kind: there is no strict cut-off point. Taken together, our results present a challenge both to traditional internalist and mainstream externalist theories of natural kind terms' reference.

Is ordinary concept of memory factive?

Vilius Dranseika

In this paper, I present new data bearing on two constraints that are often taken to be essential features of our ordinary use of ‘remembering’ and ‘having a memory’: the factivity constraint (i.e. that one can be truly said to ‘remember’ some event only if that person originally experienced or observed that event) and the strong previous awareness condition (i.e. that remembering presupposes identity between the person who remembers an event and the person who originally experienced that event). Studies were conducted in Lithuanian language (4 studies, combined N=746). To the best of my knowledge, this is the first attempt to address empirically the question whether these two constraints are features of our ordinary concept of memory. The present set of studies suggests that the factivity constraint and the strong previous awareness condition are not essential features of our ordinary use of ‘remembering’ and ‘having a memory of’. Concerning the factivity constraint, artificial memory and misidentified dream memory vignettes involved violations of factivity, and in all these cases study participants tended to agree that the agent ‘remembers’ or ‘has a memory’. The fact that study participants tended to agree that the agent ‘remembers’ and ‘has a memory’ in cases of having implanted other people’s memories, suggests that the ordinary notion of memory is not bound by the strong previous awareness condition either. When participants were given a choice to indicate that what is described in the vignette is not a case of having a memory of an event but a case of having an experience of a sort that would be experienced by someone who indeed has a memory of such an event, they were still more likely to choose an option ascribing a memory rather than an option ascribing an experience like the experience of someone who has such memory. These findings, of course, should be taken as only the first preliminary and very limited step in the direction of better understanding of constraints that rule our ordinary notion of remembering. Among limitations of this study, I would like to stress the very limited set of experimental vignettes used (misidentified dream memories, artificial memories, memory transfer), the fact that it is unclear whether the results would generalize to other languages than Lithuanian, as well as that some of the vignettes were based on science fiction scenarios. In summary, the data provided in this report provide some evidence to motivate skepticism concerning whether the factivity constraint and the strong previous awareness condition are essential features of our ordinary use of ‘remember’.

Zombie Intuitions

Eugen Fischer and Justin Sytsma

Our study contributes towards debunking the key intuition at the root of David Chalmers' 'zombie argument' (Chalmers, 1996). This influential argument against the supervenience of phenomenal consciousness on the physical relies on the intuition that 'zombies' which are physically and behaviourally like us but lack conscious experiences are logically possible. This intuition proved strikingly controversial, as some philosophers strongly shared it, while others found it utterly counter-intuitive to allow that beings with bodies and behaviours like us could lack conscious experiences (review: Kirk, 2019). We believe this clash of intuitions is due to theoretical and linguistic factors. Our study examines a linguistic factor that can be examined experimentally with lay participants. We develop and experimentally test a psycholinguistic explanation of Chalmers' intuition and briefly discuss the philosophical relevance of our findings in the light of subsequent development (Chalmers, 2002) and metaphilosophical discussion (Cappelen, 2012) of the zombie argument. We finally note general methodological consequences for developing philosophical vignettes. We draw on neo-Gricean pragmatics (Levinson, 2000; Garrett & Harnish, 2007) and the heuristic-analytic theory of reasoning (Evans, 2006) to suggest that judgments about verbally described scenarios are based on situation models shaped by automatic stereotypical inferences. These inferences are subject to a 'salience bias' documented with reading-time measurements and pupillometry (Fischer & Engelhardt, 2019; in press), for inferences from verbs. We suggest this bias extends also to inferences from nouns: When words ('zombie') have several related senses (incl. 'Hollywood' and philosophical sense), and one of these senses is clearly dominant and functional for the interpretation of less salient uses (Giora, 2003), then contextually inappropriate inferences licensed only by the dominant sense are triggered also by the less salient uses and go through to influence judgments. We hypothesise: (1) the dominant 'Hollywood' sense of 'zombie' is stereotypically associated with several features suggestive of lack of conscious experiences (rigid stare, etc.); (2) stereotypical inferences of these features are triggered by statements of the zombie argument, and influence intuitive judgments about philosophical zombies, despite being contextually defeated by mention of bodily and behavioural similarity; (3) appropriate alternative descriptions of philosophical zombies lead to different judgments. To examine our first hypothesis, we used typicality ratings to assess how strongly various features are associated with the noun 'zombie' (Rosch & Mervis, 1975). We assessed this first for individual features and then feature clusters, as identified by cluster analysis. In a second study, participants assessed how plausible it is that various beings (including zombies) and (otherwise unspecified) 'beings' with these clusters of properties should lack conscious experiences. Goodness of example-ratings elucidated participants' understanding of 'conscious experiences'. We then examined whether the zombie stereotype would still influence judgments about duplicates of humans referred to as 'zombies', even when duplicates were described as having bodies like us and behaving like us (as by Chalmers). In the first experiment, vignettes described such philosophical zombies either as 'zombies' or as 'duplicates', and participants rated attributions of typical zombie features and of lack or possession of conscious experiences. Further experiments examined the effect on such ratings of framing vignettes in third-person vs first-person terms, and of increasingly demanding attention and comprehension checks. Findings supported our hypotheses.

Much at Stake in Knowledge

Alexander Dinges and Julia Zakkou

Many philosophers have voiced intuitions to the effect that knowledge ascriptions are sensitive to stakes in the sense that we become less inclined to ascribe knowledge when it is important to be right. These intuitions, however, have resisted empirical confirmation. A number of studies fail to find any effect of stakes on knowledge ascriptions at all (Feltz and Zarpentine, 2010; Buckwalter, 2010, 2014; Buckwalter and Schaffer, 2015: 216–218; Rose et al., 2019; Francis et al., ms). Studies using the so-called evidence-seeking paradigm yield more promising results (Pinillos, 2012; Pinillos and Simpson, 2014; Buckwalter and Schaffer, 2015: 208–209; Francis et al., ms). It can be argued though that these studies confuse knowledge ascriptions with normative judgements about what the protagonist of the story should do (Buckwalter and Schaffer: 207–218; Rose et al., 2019: 240) or that they confuse stakes effects with salient alternative effects because people are independently known to think of more error-possibilities when the stakes rise (Buckwalter and Schaffer, 2015: 222). We present a novel experimental paradigm to test stakes effects on knowledge ascriptions, namely, the retraction paradigm. Participants are asked to imagine themselves in e.g. a bankcase-like situation where they assert “I know the bank will be open.” For one group of participants, the story continues such that the stakes are revealed to be high. For another group, the stakes remain low. Then participants are asked whether they would be more likely to retract their previous knowledge ascription or stand by it. They could also rate their confidence in the response they gave. It turns out that participants become substantially more inclined to retract when the stakes rise. This outcome is confirmed for a corresponding version of the familiar typo stories. We suggest that these results are less amenable to the indicated confusion between knowledge ascriptions and normative judgements about what to do. We go on to present a further study that aims to put pressure on the idea that our results confuse stakes effects with salient alternative effects. In a pretest, we collected error-possibilities for the bank cases that would naturally come to people’s minds. The basic strategy was to ask participants what could happen such that their belief that the bank will be open turns out to be false. We systematized their responses and ended up with four error-possibilities (roughly: changed opening hours, holiday, false memory, staff training). In an exact replica of the bank case study from above, we added a further screen where we presented a list of the previously collected error-possibilities and asked participants to indicate whether they had thought about them. Then we assessed whether their retraction judgements were mediated by the number of error-possibilities they had thought of. As it turns out, there is an indirect effect of stakes on retraction via the number of error-possibilities people think of. A much stronger direct effect remained though. Thus stakes effects are unlikely to result just from the fact that people think of more error-possibilities when the stakes rise.

How bias-resistant are moral philosophers?

Alex Wiegmann and Joachim Horvath

Thought experiments figure prominently in philosophy, especially in moral philosophy. Our intuitions about hypothetical moral scenarios are commonly treated as (defeasible) evidence and play an important role in developing and evaluating moral theories. However, the method of relying on casebased intuitions has received strong criticism based on empirical studies in experimental philosophy. Several studies found that moral intuitions can be influenced by factors that are apparently morally irrelevant, such as order of presentation. Based on such and similar findings, it has been argued that the method of cases is not reliable and its use in moral philosophy should be restricted. The so-called expertise defense is probably the most popular and promising reply to this challenge. It claims that findings of the described kind do not really pose a threat to the use of case-based intuitions because the respective studies exhibit an important limitation: Participants in these studies were lay people and, according to the expertise defense, it seems plausible to assume that experts' (i.e., professional moral philosophers') intuitions are largely immune to irrelevant factors. The expertise defense is an empirically testable hypothesis, but only a few studies on a rather narrow range of irrelevant factors have been conducted so far. We aim to advance the debate by presenting five cases involving robust and successfully replicated biases to expert moral philosophers and lay people, in order to compare to what extent both groups are influenced by these morally irrelevant factors. One example is the "decoy-effect" (Huber, Payne, & Puto, 1982), according to which adding a new option to a set of options can increase the probability of choosing an option that was already included in the original set. We apply this effect as follows: participants in one experimental condition are presented with a "trap-door" variant of the trolley-dilemma, where the train can be stopped by opening a trap-door in a bridge, causing a heavy person to fall on the tracks. With only two options—not to intervene or to open the remote-controlled trap-door—participants' responses are usually roughly split. This is probably due to the fact that both options exhibit features that speak for or against them: the net sum of lives saved favors opening the trap-door, whereas using the one person as a mere means speaks against it. In the other experimental condition, we add a decoy-option: a trap-door in another bridge with two persons, who would also fall on the tracks. We predict that adding this option, which is clearly worse than the original trap-door option, but better than non-intervention in terms of lives saved, will result in more positive ratings for the original trap-door option. Overall, we will test five cases involving robust biases with expert moral philosophers and lay people, in order to compare how these groups are influenced by these morally irrelevant factors, and we plan to present our results at the Experimental Philosophy Conference in Bern (we are currently in the process of completing our study).

Bioethical judgments: effects of training

Nora Heinzelmann

Judgments about issues from applied ethics are based on factual as well as moral beliefs. For instance, whether someone thinks that genome editing in humans is ethically permissible depends both on her knowledge of genetic technologies and her moral views about the sanctity of human life. This raises the question of whether and how information that may shape these prior beliefs may also affect laypeople's bioethical judgments.

We present a research project investigating to what extent, if any, individuals change their minds on bioethical questions when they are exposed to descriptive as well as ethical training. By way of example, we focus on judgments about issues surrounding novel genome editing technologies such as CRISPR/Cas9 that have recently gained much academic and public attention.

We present results from four pilot studies. The first three served to design a bioethical questionnaire as a measure for bioethical judgments. The fourth study explored the effect of training on these judgments.

Concurring with the literature, we find that bioethical judgments are highly context-dependent. On the one hand, they are determined by individual features such as gender or religiosity. For example, women are more opposed to genetic testing in animals than men. On the other hand, bioethical judgments vary with the context, purpose and kind of genetic technology. For instance, participants judge that genome editing is morally more permissible in adults than in embryos, and they regard it as less permissible when employed for enhancement rather than treatment.

In our planned presentation in Bern, we focus on the fourth pilot study. Here, we measured bioethical judgments in two groups of high school students before and after they attended a two-day workshop. During the workshop, participants learned about either personalised oncology (treatment group) or project management (control group). The workshop on personalised oncology covered topics such as genetic technologies, personalised medicine and moral reasoning.

We measured judgments about ethical and factual questions before and after the workshops as well as confidence about these judgments. We then investigated differences both between pre- and post-workshop responses and between the two groups.

Our three main results are the following. First, we find a highly significant boost in confidence for the training group. Participants in this group were more confident about their factual and ethical judgments after the workshop. In line with earlier research, we expected that confidence about factual questions would increase irrespective of the correctness of the given response. In fact, we find that the confidence boost for incorrect responses is at least as strong as it is for correct ones. None of these findings were replicated in the analysis of the control group's responses, suggesting that the confidence boost is due to specific training and not merely an effect of repetition.

Second, confidence within the ethical domain was weakly correlated with general consensus about the respective item, i.e. the degree of agreement with other

participants. As this already applies to participant's confidence before attending the workshop, it might be explained by prior awareness or expectation of consensus.

Lastly, confidence was also correlated with degree of (dis)agreement. That is, the more extreme a participant's ethical view on an issue, the higher was her confidence. This resonates with the commonplace that moral extremists have the loudest voice whilst moderate positions are less visible, but this finding needs further support from future research.

The Stability of Philosophical Intuitions: Failed Replications of Swain et al. (2008)

Adrian Ziółkowski

The paper focuses on three replication attempts of a study originally conducted by Swain, Alexander and Weinberg (2008). We not only present data that do not corroborate the original findings but also provide an in-depth discussion of factors we observed in the process that negatively affect the replicability of the original experiment. We will use this case study to illustrate the importance of precise reporting for the replicability of experimental studies. Experimental philosophers regularly report data that seem to be at odds with philosophical consensus. Data on discrepancies between philosophers' and folk intuitions provide the crucial premise for the argumentation against the use of intuitions in philosophy put forward by the representatives of negative experimental philosophy (also referred to as "experimental restrictionism"). Since certain intuitions are not universal, as experimental restrictionists argue, the consensus among philosophers should not be used as evidence in favor or against philosophical theses. Another argumentative strategy leading to a similar conclusion points at instability of intuitions, because, if intuitions are shaky, they are not reliable and should not be trusted. In their widely cited article, *The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp*, Swain et al. (2008) use the latter strategy to argue against the use of intuitions in epistemology. They report data that, purportedly, demonstrates instability of folk epistemic intuitions regarding the famous Truetemp case authored by Keith Lehrer. What they found is a typical example of priming, where presenting one stimulus before presenting another stimulus affects the way the latter is perceived or evaluated. In their experiment, laypersons were less likely to attribute knowledge in the Truetemp case when they first read a scenario describing a clear case of knowledge and more likely to ascribe knowledge when they first read a vignette describing a clear case of nonknowledge, with subjects that were not primed falling somewhere in the middle. We tried to replicate Swain et al. findings in three experiments: two conducted in English and one devised in a different Indo-European language. We found no priming effect for knowledge ratings regarding the Truetemp case – laypersons were similarly likely to attribute knowledge in all three investigated conditions (primed with a clear case of knowledge, primed with a clear case of nonknowledge, and not primed). We conducted a meta-analysis of the data collected in our three replication attempts separately for the positive and negative priming effect. In both cases, the combined effect size confidence interval indicates that the effect in question either does not exist, or – if it does – is very small, which puts its theoretical importance in question. Our three failed replication attempts of Swain et al. (2008) are not decisive as to whether the priming effect in question occurs. Nevertheless, the collected data puts Swain et al. conclusions about instability of epistemic intuitions in jeopardy and shift the burden of proof on them – if they want to argue that epistemic intuitions are unreliable due to instability, they first need to provide more evidence that epistemic intuitions really are unstable.

How Young Minds See Future Minds: Uncanny Valley and Ascription of Different Types of Attributes to Robots

Robin Kopecký and Michaela Košová

The question of the impact of new technologies on human society becomes ever more pressing. Children of next generations will be born into the world full of new stimuli that will form their thinking in the crucial years of development. Nowadays interactions between child and machines are being studied on both philosophical and technical levels (Tung 2011, Belpaeme et. al 2013). It is possible that children will ascribe certain specifically human traits to robots that will exhibit certain behavioural patterns. In our interview study with children and teenagers (N=209, 109 F, age range 6-17, mean age=11.05) we decided to test the children's reaction to different types of robots. In the first part of the interview, the children were shown pictures of 6 different robot faces created by an industrial designer. The pictures were intended to be a spectrum from the most robotic to the most human-like face. The children were supposed to judge how much the particular face appears friendly to them on a 6-point Likert scale. As we hypothesised, in the resulting graph we observed an obvious uncanny valley. The friendliest robot was the one with eyes, but without other facial features, while the least friendly was the robot with human facial features set on the metal head, together with two robot heads with no facial features whatsoever. The robot that had all the facial features, hair and skin, scored as the second most friendly-looking robot. In the second part of the interview, the children were asked to judge whether they would ascribe certain attributes (sense perception, thinking, emotions, self-reflexion, life, soul and freedom) to different beings or entities described to them. The training examples were a dog and a child. Then they were asked about a robot with basic perception and motor abilities that can also do complicated maths and talk, and a robot who is also capable of being an active participant of a dialogue and can talk about how it perceives the world in a way indistinguishable from a real person. The results showed that children were more prone to ascribe certain human abilities to the android than to the mathematical robot (esp. emotions and freedom). Both robots scored highly in sense perception and thinking, and android scored also in self-reflexion. However, even the android didn't score when it comes to the ascription of life and the soul. Participants ascribed soul and life only to the living beings. The results show that the children are capable to ascribe certain level of cognition, emotions and even freedom to sufficiently complex robots, but are still reluctant to see robots as living beings with a soul. The concept of the soul may be understood as an essential moral core of the person that is capable of moral conduct and deep interpersonal bonds and which is conceptualized differently than the mind (Richert & Harris 2006, 2008). Despite their behavioural complexity and intelligence, robots might not manage to cross the boundary between machines and full-fledged persons in the eyes of children.

Ideas are cheap: When and why adults value labor over ideas

Pascal Burgmer, Matthias Forstmann and Olga Stavrova

“Ideas are cheap. Execution is everything.” Many entrepreneurs might agree with this statement by Scott Adams, cartoonist and creator of the Dilbert comic strip. However, ideas are highly appreciated and even protected by intellectual property and patent law. But do lay people value ideas over labor or do they rather share Scott Adams’ perspective? This question has received very little empirical attention so far. Yet, whether or not people value ideas over their execution matters for how they ascribe ownership, authorship, and other material as well as non-material rewards and punishments for collaborative work.

The empirical evidence in this domain is mixed. People assign ownership on the basis of perceived labor—particularly when such labor adds to the value of a particular piece of work. Additionally, research on the “effort-heuristic” indicates that people value work (e.g., a painting) more if they believe it to reflect a lot of effort. Recent developmental research, however, suggests that children at the age of 6 begin to value ideas over labor. Specifically, 6 year olds were found to prefer a picture containing their idea over a picture that they merely made based on someone else’s idea. Furthermore, they also awarded ownership of the picture to the idea giver rather than the laborer, whereas 4 year olds were indifferent to this distinction. Thus, people do value ideas, but they also do not entirely dismiss labor when estimating a work’s value or when resolving ownership issues.

In the current paper, we report a total of seven studies ($N = 1,463$) to clarify when and why adults might value the one over the other. In these studies, participants usually saw one or more short scenarios depicting both an idea giver and a laborer who collaborated to create something (e.g., a computer application, a picture, or a meal). Participants were then asked about their intuitions about who deserved ownership and monetary compensation for the creation. Studies 1a–1c found that participants valued the contribution of the laborer more than the contribution of the idea giver. This labor-valuation effect emerged even when participants themselves were idea givers (Study 1b), and it was replicated across different populations (including legal professionals, Study 1c) and contexts (e.g., art works and businesses, Study 2). Studies 3a and 3b established perceived effort as a central psychological process behind the labor-valuation effect. Finally, Study 4 extended the effect to the realm of praise and blame judgments, showing that laborers receive more praise for positive outcomes, but less blame for negative outcomes, relative to idea givers.

Taken together, our studies suggest that adults value labor over ideas, because they perceive contributing labor (vs. ideas) as more effortful. This effect is moderated by the valence of the mutual creation, suggesting that people have different intuitions when a creation turns out to be a failure (vs. success). These findings have implications for psychological and philosophical considerations about perceptions of agents who collaborate to create things.

The norm effect in causal selection - Empirically investigating the role of the relation between norm and outcome

Sophie Lusser and Alex Wiegmann

Researchers in both Psychology and Philosophy have been examining judgments of actual causation, the question to which extent a certain action (or event) caused a given outcome. One interesting phenomenon occurs for situations with the following features: two actions are necessary produce a certain outcome (conjunctive causal structure), and only one action complies with a norm at place, while the other action violates it. In such situations, people tend to ascribe more causality to the norm-violating action and also choose it over the normcomplying action if asked to select the cause that brought about a certain outcome. Several competing explanations have been proposed for this phenomenon called the “norm-effect”. In the current paper, we investigated a previously neglected factor that can shed some light on what is driving the norm-effect, namely the relation between the norm and the outcome. First, a norm might be introduced with the aim to produce/prevent a certain outcome (henceforth: goal-orientedness) or a norm might be accidentally in place. Second, following a norm might ensure that a certain outcome is prevented/produced or it does not (henceforth: effectiveness). In our experiment, each of N=400 participants read two vignettes with a conjunctive structure in which two agents performed similar actions. In each vignette, one agent’s action was normconforming, while the other violated a norm. We manipulated the goal-orientedness (present vs. absent) and effectiveness (present vs. absent) of the norm in a 2 x 2 between-subjects design. For instance, in the Bridge case two drivers cross a bridge from opposite sides, causing the bridge to collapse due to the combined weight of the two cars. In the effective norm conditions, a traffic light signal prevents the negative outcome by allowing only one car at a time to cross the bridge. In the ineffective conditions, the parents of one driver do not allow their child to cross the bridge, but there was no norm in place that in general—and not only in this specific situation—would rule out two cars crossing the bridge at the same time. In the goal-oriented conditions, the norm was specifically introduced to prevent the collapse of the bridge, in the non-goal-oriented conditions the norm was introduced for other reasons (e.g., to prevent damage of side mirrors because the bridge is narrow). After reading either vignette, each participant was asked to select either the norm-conforming agent’s action, the normviolating agent’s action, or both as the cause of the negative outcome. We found the same result pattern for both vignettes. While goal-orientedness didn’t affect participants’ responses ($p = .64 / p = .75$), effectiveness had a strong effect ($p < .001 / p < .001$)—and the two factors did not interact. More specifically, in conditions in which following the norm ensured the prevention of the negative outcome, the norm-violating agent’s action was chosen by one-third of participants as the cause, while in non-effective scenarios his action was almost never chosen. We discuss the implications of our findings for several theories proposed to explain the norm-effect.

Moral responsibility: The attribution of blame is influenced by culturally universal features and culture-dependent aspects of causal involvement and social roles

Albert Newen, Pascale Willemsen and Kai Kaspar

What are the main features that influence our attribution of moral responsibility and what is the role of culture? It is well known that intentionality as well as the valence of the outcome strongly influence our evaluation of moral responsibility. To investigate the relevance of social hierarchy and causal responsibility, we used a vignette introducing a joint company activity of manager and technician. In our scenarios we focused on the attribution of blame for the negative sideeffect of the joint action and asked for the attribution of blame to both, a manager/boss and a technician/employee. At the same time we kept the intentionality of manager and technician concerning the side-effect neutral and constant. Interestingly, social hierarchy strongly modulates the attribution of blame. Furthermore, we put this into an intercultural context comparing completely new data from Germany, United Arab Emirates (UAE), Poland, USA and China. In all five cultures the main effect of hierarchical difference between boss and employee has the same tendency and is strongly significant: the boss deserves more blame than the employee. Furthermore, the experiment was designed to clarify whether social hierarchy is a new factor, independent from causal responsibility, to be accounted for in a systematic theory of attributing moral responsibility. To prove this we designed the vignette such that we distinguish the role of the decision maker from the social authority. While in our former studies we arranged it such that the boss is not only the social authority but at the same time makes the decision, we now produce a scenario of a technician in a company who is enabled to make one specific and important decision while the manager remains the ultimate authority. Thus we have a normal decision making scenario (manager makes the decision) and a reverse decision making scenario (technician makes the decision). Furthermore, we ask about the expectation to care about the negative side effect being in a specific social role, i.e. manager or technician. This disentangles causal responsibility and social role. We predicted that in the reverse decision making scenario the boss still receives more blame than the employee even if the causal responsibility of both is equal; either this is true for all cultures or dependent on the difference in the understanding of social hierarchy. The results so far: The contrast between the normal and reverse decision making scenarios has a very clear effect for all cultures concerning the causal responsibility: in the normal decision making condition there is a significant difference between a high level of causal responsibility of the manager and a low level of the technician. In the reverse decision making condition this difference is gone for all cultures, i.e. the level of causal responsibility for manager and technician is basically the same. Now, what about the level of blame attribution? Our prediction that the manager still receives more blame in the reverse condition is confirmed for China and the USA, there is a trend observable for UAE and Germany; only Poland does not show this effect. What about the intercultural differences? We aim to explain these differences by including our additional data: (i) expectation to care about the negative side effect for manager and technician (social role) and (ii) cultural data of the test persons (standard test of horizontal versus vertical individualism and horizontal versus vertical collectivism), i.e. social hierarchy understanding. The result of this intercultural study is, first, to establish social role as a strongly relevant factor modulating the attribution of moral responsibility and to work out the aspects of social role which are interculturally similar and those which are culturally specific.

Pragmatics, levels of meaning, and speaker commitment

Alison Hall and Diana Mazzarella

What a speaker explicitly communicates, as opposed to implicates, is held by Contextualists (e.g. Sperber and Wilson 1995, Carston 2002, Recanati 2004) to go beyond Grice's notion of 'what is said', which is equated with the utterance's linguistically encoded meaning after reference assignment and disambiguation. This pragmatically enhanced level of meaning is labelled 'explicature'. 'Intrusion' of non-linguistically mandated pragmatic processes raises the question of how to distinguish explicature from implicature.

Borg (2017), Weissman and Terkourafi (2018), and others, argue that the lying misleading distinction tracks the explicit-implicit distinction better than traditional truth value judgment tasks do: roughly, explicitly communicating something false with the intention to deceive is lying, whereas implicating it is misleading. Borg (2016, 2017) claim lying-versus-misleading judgments favour the more minimal what is said over explicature as the correct notion of explicit content.

We argue that the lying-misleading test is flawed: the process of judging whether the speaker lied does not uncover communicated content, but instead encourages the judge to extract the minimal content recoverable by linguistic decoding (plus reference assignment and disambiguation), because this minimal content, being largely free of pragmatic inference, is what the judge is in a position to indisputably hold the speaker responsible for.

We further argue for a more indirect method of gauging what hearers entertain as explicit content. Using uncontroversial examples of asserting (i.e. cases without pragmatic intrusion), presupposing, and implicating, Mazzarella et al (2018) measured reputational costs incurred by speakers who deliberately conveyed something false, finding that, *ceteris paribus*, people are "significantly more likely to selectively trust the speaker who implicated *p* than the speaker who asserted or presupposed *p*". We present results of a study currently in progress to compare explicature to what is said and implicatures. Examples include (1), adapted from Simons (2017):

1. What are you getting your mother for her birthday?
 - a. I'll buy some flowers to give her
 - b. I'll buy some flowers
 - c. She likes flowers

Reply (c) implicates the speaker will buy flowers and give them to her mother; in (b), this is part of explicature. We predict that pragmatically inferred, explicated material, as in (b), is treated by participants as invoking a similar level of speaker commitment as is characteristic of content that is encoded, or involves only reference assignment and disambiguation, as in (a). We discuss the implications for the debate about what levels of meaning feature in the comprehension process.

The Aesthetic Self

Joerg Fingerhut

Aesthetic values and our engagements with the arts are an essential part of our identity. The aesthetic preferences we have, make us the person we are. We present a series of empirical findings of an Aesthetic Self-Effect supporting such a claim. Counterfactual changes in aesthetic preferences - from liking pop to liking classical music - are perceived as threatening a person's identity. The effect is as strong as the one found for moral changes - such as altering political partisanship or religious orientation - and significantly stronger than changes in other categories of taste, such as food preferences (study 1, N = 251, 359). Using a multidimensional scaling technique to compare perceived aesthetic similarities amongst musical genres, we determined that the choice of our specific example genres - i.e., pop and classical music - was not responsible for the magnitude of the Aesthetic Self-Effect. We found furthermore that neither directionality, the central-peripheral dimension, or movement between areas on the similarity map of our genres had an impact on our self-measure, but only exceeding an aesthetic threshold (study 2, N = 45, 364). Further studies support a generalized Aesthetic Self-Effect beyond the musical domain: e.g., general changes in art preferences from more traditional to abstract art also elicited a strong Self-Effect (study 3, N = 237). By exploring the breadth of this effect we found evidence of an additional Anaesthetic Self-Effect. That is, scenarios that describe the original adoption of an aesthetic preferences - for instance from not caring to caring about music, art or beauty - also elicited strong judgments from our participants. This effect is significantly stronger for the aforementioned classically aesthetic fields compared to adopting other leisure activities, such as hiking or playing video games (study 4, N = 305). This is evidence for genuine aesthetic self: both, our taste in music and the arts, as well as our general ability to appreciate the arts and beauty belongs to the most central features of our identity.

There has been some recent evidence of the central role moral values have for our identity, that prompted researchers to proclaim a "moral self" (Strohminger and Nichols 2014; Prinz and Nichols 2016) Our paper builds upon this research and extends it systematically towards the other philosophical domain of values: aesthetics. We are not just moral selves but crucially aesthetic selves as well. The main aim of this paper is, therefore, to explore whether taste changes actually exert an impact on perceived identity. In a series of experiments, we presented participants with vignettes that prompt them to imagine a taste change for themselves and then asking them to what extent they perceive themselves to be the same person after the change. The striking finding across all our experiments is that taste changes are among the changes that present the biggest threat to the identity of a person.

Let's lay it on thick!
**On the relevance of thick ethical terms for metaethics and moral
psychology**

Pascale Willemsen & Kevin Reuter

Philosophers and linguists usually distinguish two types of evaluative terms and concepts, namely “thin” and “thick” ones. Thin terms evaluate an object as, for instance, right, permissible or wrong, yet they don’t explicate in which way the object is right or wrong. Thick ethical terms and concepts provide such information in addition to evaluating the object. Typical examples are “rude”, “courageous”, “trustworthy”, or “dogmatic”. Describing an agent as courageous makes a descriptive judgment that the agent is willing to take risks, and it evaluates taking risks positively – in contrast to being reckless. Thick ethical terms and concepts have played a major role in the most fundamental metaethical debates, such as the debate between cognitivists and non-cognitivists. Both sides make far-reaching assumptions about how thick terms and concepts work in ordinary language, and about whether it is possible to disentangle the evaluative from the descriptive dimension. However, there is no empirical evidence as to how and for what we use thick ethical terms and concepts. In this talk, we will present empirical evidence from three experiments. We will discuss to what extent this evidence supports the Disentanglement Argument and, thus, can be used to argue against Non-Cognitivism. In doing so, we will put emphasis on the severe methodological challenges that come with research on thick concepts. We will suggest that no matter how severe the challenges, thick ethical terms and concepts need to be addressed empirically for their outstanding role in metaethics but also moral psychology.

Do intentions really matter? An experimental investigation of demonstrative reference

Wojciech Rostworowski, Bartosz Maćkiewicz and Katarzyna Kuś

One of the controversies in the philosophical debates on the nature of reference concerns the demonstrative reference: namely, the relation which holds between a demonstrative expression (a simple: "this" or "that"; complex: "this table", "that black dog" etc.) and an object. It seems reasonable to think that the semantic value of a demonstrative is the object which the speaker intends to be that value and which, optionally, satisfies some further conditions (e.g., see: Åkerman 2015, King 2014, Michaelson 2013, Speaks 2016, Stokke 2010). Let us call this view "Intentionalism". An alternative for Intentionalism are the accounts on which reference is determined by some aspects of the "external" context, namely, by a broadly construed salience in the sense that the value is the most salient object in the context (see: Gauker 2008, Rostworowski & Pietrulewicz 2018). Although the debate about demonstrative reference is active for some time, there is little empirical evidence on that matter in the philosophical discussions. This is surprising, given that many arguments directly appeal to linguistic intuitions and thus call for empirical validation. In our study we presented the participants with short scenarios describing situations where a speaker made a statement of the form "This F is G". There were always two different objects - one determined by the speaker's intention, the second one determined by an alternative factor (like the descriptive content in the nominal "F", or the speaker's gesture) - also differentiated in terms of satisfying predicate G. Our overall prediction was that in the scenarios where the object determined by a non-intention factor had the ascribed property, the speaker's statement will be judged as true (contra Intentionalism); accordingly, it will be judged as false when the object determined by the non-intention factor did not have the ascribed property. Our study provides counterevidence to Intentionalism. It was observed that the object denoted by the nominal or an act of demonstration is preferred as the referent to the object which is merely intended by the speaker. We will argue that these results cannot be easily explained by introducing competing intentions (cf. Speaks 2016) or by the distinction between speaker's reference and the semantic reference. Finally, we outline a theory of demonstrative reference, on which the semantic value of a demonstrative is determined as the object best satisfying various "accessibility" criteria (cf. Gauker 2008).

Aesthetic judgments as relativised to specific circumstances of judgment

James Andow

Experimental philosophy of aesthetics has explored what extent the folk are committed to realism about aesthetic judgments. Extant work has focused on the extent to which participants are inclined to reject normativism. Normativism is a key commitment of realist positions in aesthetics and is the claim that aesthetic judgments and statements have correctness conditions, which do not vary between subjects, such that there is a fact of the matter in cases of aesthetic disagreement.

The emerging picture is that the folk strongly and almost universally reject normativism and that thus there is no strong realist tendency in folk thinking about the aesthetic. This emerging picture has thus been taken to dissolve the traditional philosophical puzzle in aesthetics of how to best account for the fact that aesthetic judgments seem to us to have intersubjective validity in light of the apparent subjectivity of aesthetic experience. This paper presents a series of studies which further enrich our understanding of ordinary thinking about the aesthetic. The results suggest that ordinary thinking about the aesthetic is not quite so vehement in its rejection of normativism.

The results also suggest that the interpretation of the results of previous studies should be revised in other ways. Previous results suggested that, at least in many cultures, the dominant trend is to reject the idea that aesthetic judgments have correctness conditions at all. The current results suggest the previous results should not be interpreted in this way: participants do think it is possible for aesthetic judgments to be incorrect but that the correctness conditions are very finely relativised to very specific circumstances of judgment.

Utilitarianism beyond sacrificial dilemmas: testing the impact of three ways of manipulating reflexivity on utilitarian judgments across six topics

Florian Cova, François Jaquet

Joshua Greene's famous and popular 'dual-process' model of moral judgment claims that typically deontological judgments are supported by automatic, emotional processes, while typically utilitarian judgments are supported by controlled, reflective processes. Extensive evidence has been gathered in favour of this claim, the most direct of which are studies showing that manipulating participants' reflexivity (using cognitive load, time pressure, or counter-intuitive reasoning problems) actually influences the rate of participants' utilitarian judgments.

However, most of these evidences have been gathered through studies focusing on one very particular instance of opposition between deontology and utilitarianism: sacrificial dilemmas (i.e. dilemmas in which one or more persons must be sacrificed to save even more persons). Despite these limitations, Greene argues that his model extends to other cases of deontology/utilitarianism conflicts, such as the existence of harmless moral violations, the moral relevance of the action/omission distinction, the extent to which ethics can be demanding, or the reasons that warrant and justify punishment.

But do findings about sacrificial dilemmas really extend and generalize to these topics? To find out, we first created a new battery of 60 scenarios (6 different topics, 10 scenarios per topic) and had it validated by professional ethicists. Then, we tested the impact of three different experimental manipulations on these different scenarios (a cognitive load, a time pressure, and a prior exposure to counter-intuitive reasoning problems). Though adequately powered, these studies failed to find an effect of these manipulations for most areas of disagreement between deontology and utilitarianism. We argue that our results shed doubt on the generalizability of previous results to moral debates outside the particular case of sacrificial dilemmas. Moreover, our results also question the replicability of earlier results about sacrificial dilemmas.

Folk-concepts of happiness and well-being: Empirical Evidence

Markus Kneer

The project examines the folk concepts of happiness and well-being. A series of experiments (total N>2000) explores to what extent ascriptions of happiness and well-being are sensitive to internal factors (psychological states) and external factors (life conditions). The results suggest, expectedly, that (i) folk concepts of well-being are more sensitive to external factors than folk concepts of happiness. Moreover, they reveal, unexpectedly, that (ii) psychological states are the dominant factor for all tested attributions, including well-being and the good life, and (iii) that the expression “happy” is ambiguous and thus that there is a plurality of folk-concepts of happiness.

Motivational Constraints and Freedom of Choice

Lukas Huber and Kevin Reuter

Adults have the intuition that desires do not impair their freedom of choice. While most four and five-year-olds do not share this intuition, six-year-olds seem to agree with adults in thinking that they could have done otherwise than following their desire. It has been proposed that six-year old children have formed a concept of free will according to which agents can cause actions (Nichols, 2005) and conceive of free will as a causal force that mediates between desires and action (Kushnir et al., 2015). In this paper, we report a new study that is the first detailed investigation examining the development of intuitions about motivational constraints in 4- to 6-year-old children. Our study shows that six-year old children only endorse freedom of choice when alternative desires are made salient. The results suggest that six-year-old children do not entertain a causal force concept a la Kushnir but rather a *conditional* concept of free will, according to which they would have acted otherwise had they followed another of their desires.

I'd feel guilty but I shouldn't: An asymmetry between prescriptive reasoning and affective experience in autonomous car accidents

Fernando Aguiar, Pilar Aguilar and Ivar Hannikainen

In recent years, autonomous or driverless cars have taken to the streets in many different cities around the world. These "moral machines" (Awad et al 2018a) need to be programmed to decide who to save or sacrifice when an accident can be foreseen, but not avoided—a question that has aroused the interest of moral philosophers. Should these vehicles risk the passenger's death in an attempt to save a larger number of pedestrians' lives, for example? Will people accept these utilitarian, driverless cars? In interesting experimental work emerging from the trolley problem tradition, Bonnefon et al (2016) conclude that people, in fact, are quite comfortable with the existence of utilitarian, autonomous cars—although with certain, important caveats (Awad 2018a: 60). In addition, people are willing to attribute blame and causal responsibility to autonomous vehicles in case of a crash (Awad et al 2018b). In our work, we ask a complementary question about passenger responsibility: Do the passengers of an autonomous car feel guilty and/or responsible for accidents beyond their control? To answer this question, we devised a 2 (car: autonomous, driver) x 2 (point-of-view: 2nd person, 3rd person) x 2 (guilt: prescribed, perceived) between-subjects design, and recruited 200 participants to take part in our experiment. In every condition, we described an inevitable accident involving a pedestrian's death narrated either in the 2nd or 3rd person. Participants were asked whether the passenger (Autonomous Car condition) or driver (Driver condition) was responsible for the accident. Additionally, participants were asked one of two questions: (1) whether the driver/passenger would feel guilty in those circumstances, or (2) whether the driver/passenger should feel guilty in those circumstances. As expected, passengers in autonomous cars were seen as less responsible than drivers. In a mediation analysis, responsibility ascriptions mediated the effect of vehicle-type on prescribed guilt: i.e., passengers' reduced responsibility entailed a corresponding reduction in the guilt that they should feel concerning the accident. The corresponding mediation analysis with perceived guilt, however, was not significant: In other words, the difference in responsibility for accidents in autonomous versus human-driven cars did not statistically mediate the difference in guilt feelings that would arise. These results reveal an important discontinuity between participants' prescriptive reasoning about moral emotions, and their actual or imagined affective experience. Simply put, laypeople reason that, because passengers are not in control of autonomous vehicles, they would not be responsible, and should not feel guilty, if an accident were to occur. This prescriptive line of reasoning, however, does not protect people from experiencing unwanted guilt in those same circumstances—in line with a "no-judgmentalist view of guilt that makes out the subjectively guilty agent as feeling as if he were morally responsible" (Greenspan 1992: 287). Our findings illustrate a particular ethical issue emerging at the humantechnology interface, which we argue deserves greater attention going forward.

What We Don't Want From A Theory Of Happiness

Andre Bilbrough

The systematic measurement of subjective happiness is receiving unprecedented attention across, in some form or another, all of the social sciences - and provides a valuable new source of information in decision-making for society and public policy.

There is however no consensus on what happiness is and how to measure it, rendering any conclusions premised on it as deeply contentious. It has been claimed, for example, that reducing unemployment has been undervalued as a policy goal relative to reducing inflation, based on their relative effects on happiness, and yet, when using a slightly different measure happiness and unemployment are actually positively associated. Similarly, the level of income at which increases in income cease to be associated with increases in happiness varies by an order of magnitude depending on the measure of happiness used.

This leads naturally to two questions when using happiness in evaluating policy: (1) what, broadly speaking, should we expect a conception of happiness to do? and (2) given such expectations, how should we measure it?

Regarding the former question, I argue that in this domain it is important that directly describing the lay conception of happiness takes priority over conceptual revision in line with philosophical and theoretical values.

The revisionist project is typically motivated around two families of desideratum: (1) explanatory power and (2) theoretical utility. Haybron, for example, suggest that the "deep causal power" of happiness should be taken as defining. This however is out of line with empirical research which indicates that (i) happiness may actually motivate much less than we'd expect and that (ii) the extent to which individuals consciously pursue happiness varies dramatically between individuals and may be much less significant overall than such a desideratum would presume. Theoretical utility, though more commonly used, seems even more problematic as a desideratum because by straying from ordinary usage towards some ideal theoretical role, we lose the ability to justifiably use large swathes of empirical research and, even more dramatically, any appeals to intuition. Without these, the value of happiness is dramatically under-determined.

Regarding the latter question, to the extent that happiness should be taken seriously as a policy objective, too little emphasis has been given to measurement of it as momentary experience, and too much (comparatively) to it as a more general life evaluation. Though, the financial and practical reasons which have motivated this are understandable, this is problematic.

Happiness as valued by society does take the momentary experience as deeply important, and there is some evidence that it may even be considered more important. More crucially, based on previous research and presenting the results of a suggestive pilot study, I argue that measuring happiness at the larger scale as a general life-evaluation is already a normatively-laden concept, in such a way that it provides less novel information as a measurement and is problematic to integrate into normative decision-making.

Moral judgments about dilemmas in the context of war – Talk in the Field of Causal and Moral Judgment

Juan Carlos Marulanda Hernández, Alex Wiegmann and Michael R. Waldmann

In the last few decades, moral psychologists have explored people's moral judgments about dilemma situations in which doing something good involves causing harm. An often investigated example is the trolley problem in which saving a group of people from being run over can only be accomplished if fewer people are harmed. Traditionally, these kinds of dilemmas have been part of the philosophical debate between deontologists and consequentialists about whether the moral status of an act depends solely on its consequences or also on features of the act itself, such as its causal status and the violation of certain rights. Some authors have appealed to the doctrine of double effect (DDE) as a way to justify causing harms in certain situations within the framework of a deontological view. The DDE states that it is permissible to cause harm as a merely foreseen side-effect of an action aimed at achieving a greater good, although harming a person would be impermissible if it is intended as a means to achieve the same end. Empirical evidence has shown that the patterns of judgment of lay people are consistent with the DDE.

The DDE has been often invoked in the theoretical discussion about ethics of war, in particular, to justify the involuntary killing of civilians during military operations where it is too costly or difficult to guarantee their immunity. Despite an extensive theoretical debate on the subject, little experimental research has so far been done on whether people's moral intuitions about dilemmas arising in warfare are sensitive to the means/side-effect distinction. Furthermore, we were interested in investigating whether the moral distinction between killing as a means versus a side effect, which is well established for dilemmas arising in peace times, is also invoked in war. We hypothesized that this distinction is also morally relevant when civilians are the victims but suspected that it is less relevant for combatants. Combatants commit themselves to being harmed so that it should be less relevant how they are harmed. To empirically test our hypotheses, we have conducted three experiments on hypothetical war dilemmas inspired by the methodology of the trolley problem. Participants were presented with different scenarios in which they had to assess the moral permissibility of saving a larger group of people through an act that kills a smaller group, either as a means or as a side effect. Across the different scenarios, the group of victims was made up of soldiers, farmers, or children. Our general finding is that in times of war people judge the damage caused to civilians as a means to achieve a good end as less permissible

than the damage caused as a foreseen side effect. However, when the possible victims were combatants, moral judgments were less sensitive to the distinction between means and side effects. In these cases, participants focused more on the outcomes than on the causal setup.

How scientists solve moral dilemmas

Petr Jedlička

The scientific arena is fertile ground for various moral dilemmas that stem either from scientific life itself (gender or other biases, misconduct) or its impact on the broader society (progress), which contributes to a lively discussion in the scientific community (New scientist, 2017). While experimental philosophy has typically investigated moral judgements in various groups and settings, one of its branches – the x-phi of science – has so far focused more on the study of concepts (Stotz et al 2004, Linquist et al 2011) than on moral issues (Machery 2016), although dilemmas have been already routinely employed to gain insights into decision making in number of scientific disciplines, such as medicine (Evans et al. 2015).

We try to remedy this partial neglect with our study on moral judgements in contemporary science. This x-phi study is part of a broader empirical research project centered on objectivity in the natural sciences, in which an interdisciplinary team of researchers (philosophers of science, sociologists, and natural scientists) explores the notion of objectivity and its various dimensions among Czech scientists (Project “Objectivity – An Experimental Approach To the Traditional Philosophical Question” sponsored by the National Science Foundation, 2018-2020)

Our project encompasses multiple approaches to objectivity with methods including qualitative (interviews), quantitative (questionnaires) and experimental (lab experiments). The x-phi part of the study was based on dilemmas gleaned from interviews with scientists (43 interviews carried out in 2018). Some of the dilemmas were inspired by real-life situations encountered by the scientists, others originated in some current controversial events or trends in science or academia (of both local and global relevance), some of which were also publicized in the media. Subsequently, these situations were transformed into generalized dilemmas. Six dilemmas in our survey typically include a conflict between a scientific virtue pitched against a scientific or human virtue (value).

They center on issues such as:

- Scientific integrity (fighting questionable practices versus obedience to authority, individual and group responsibility for scientific fraud)
- Stereotypes in science (gender or ethnic equality versus academic freedom and autonomy, clash of Eastern and Western epistemic traditions)
- Ethical questions (scientific progress versus ethical concerns in human medical research)

In April 2019, these dilemmas were presented to a group of working natural scientists from leading Czech research and academic institutions in an online questionnaire (currently N= 442).

In my paper, I want to demonstrate how various groups of scientists – based on their demographic characteristics (sex, age, education etc.), disciplines (physicists, biologists etc.), position in the institutional hierarchy or accomplishments – handle the dilemmas and how sensitive they are to particular moral questions.

Human Soul behind the Face of the Future

Michaela Košová and Robin Kopecký

Recent debates about the problem of human enhancement eventually have to address the question of personal identity (Schneider, 2008). The aim of this paper is to draw attention to a different aspect of personal identity that has become more important with the rise of experimental approach towards philosophy. We believe that it should play a crucial role in the enhancement debate, together with obvious metaphysical problems that will not be addressed here. Firstly, various x-phi and psychological studies point to the fact that the folk concept of personal identity is socially determined (Prinz & Nichols, 2016) and that the preservation of the true self of a person stands and falls with the positive moral traits that figure strongly in interpersonal relationships (Strohmingner & Nichols, 2014, 2015; Strohmingner et al., 2017; Heiphetz et al., 2017, 2018; Tobia, 2015, 2016). Secondly, despite the general knowledge about the role of the brain, many studies show that people cannot entirely free themselves from dualistic thinking (Bering & Bjorklund, 2004; Bloom, 2004; Mudrik & Maoz, 2015). Both children, adults and even people from certain distant cultures view the brain and the mind as connected to intellectual capacities and perishable in time, while soul is considered a constant entity connected to moral conduct, interpersonality, survival after death, etc. (Richert & Harris, 2006, 2008). We will introduce our own research to support the previous findings and to show that the “essential moral self” and the concept of soul both capture the same aspect of a person. We will briefly mention the results of our interview study with children about personal identity that supports the “essential moral self” hypothesis, but the main focus of the talk will be an online questionnaire study on dualistic intuitions with nearly 3000 adult respondents. The results show that the respondents tended to ascribe to the brain competences connected to intellectual and bodily operations, memory and sense perception, while ascribing significantly lesser role to the brain when it comes to moral, emotional and deeply interpersonal competences and preservation of personal identity. In the case of the soul, the pattern of answers was exactly the opposite. Based on these findings we conclude that the “essential moral self” shares certain crucial characteristics with the folk concept of the soul: both concepts are closely connected to moral and interpersonal traits, they are both seen as facilitating the preservation of personal identity, and physical or purely intellectual traits are not usually associated with them. When it comes to the debates about the future of the human nature, physical or intellectual enhancement might be less problematic than moral or emotional one. Considering the intuitions of people living today and what they care for the most, we should strive to see all the possible enhancements through their ability to preserve or disrupt the true human self - the “soul” – an inner person capable of moral conduct and forming deep interpersonal relationships with other moral agents.

Towards a psychologically realistic understand of “could have done otherwise”

Marianna Leventi and Pascale Willemsen

In the classical philosophical debate about free will and determinism, one principle is of crucial relevance, namely the Principle of Alternative Possibilities (PAP). PAP states that an agent can only be free and morally responsible, if she could have done otherwise. If an agent's actions were fully determined by antecedent conditions beyond her control, the agent cannot be considered free or morally responsible. Incompatibilists understand PAP in a metaphysical sense, meaning that at the time of the agent's decision, it was open whether the world would continue in this or that way. As a consequence, determinism is claimed to be incompatible with free will and moral responsibility, as it entails that at any point, there is only one way. On the other hand, compatibilists tend to understand PAP in an epistemic way. Given that the agent had certain reasons to do A, it would be absurd to claim that it was perfectly open whether she would choose A or B. But had she had other information, she might have decided otherwise. Such an understanding of PAP makes determinism compatible with free will and moral responsibility. Experimental philosophers so far have greatly relied on a metaphysical understanding of PAP (Bear and Knobe, 2015; Nahmias et al., 2006; Roskies & Nichols, 2015) and tested whether the folk make judgments in line with such a metaphysical understanding. If they ascribed free will in the absence of alternative possibilities, they concluded that the folk were compatibilists. We argue that both philosophers and experimental philosophers have dedicated their time and efforts to the wrong kinds of questions. There is a growing body of evidence that people do not understand PAP or determinism in the way philosophers do (Nahmias, Morris, Nadelhoffer, & Turner, 2005; Nichols and Knobe, 2007). Further, there is no evidence that even understanding these very technical and abstract concepts matter for their everyday life. Nevertheless, it seems that in our ordinary life, “not being able to do otherwise” does matter – not in a metaphysical or epistemic but a very pragmatic sense. We typically do understand duress or coercion, obsessive-compulsive disorders, or addictions to limit an agent's free will and moral responsibility. The legal system takes factors like that into account as well. However, we still lack a systematic, controlled investigation of the relevant factors. In this talk, we will present three pre-registered experiments that deliver the first evidence of this kind. Taking inspiration from Woolfolk et al. (2008), we presented participants with scenarios in which an agent is forced to kill another person. We manipulated a) how strong this force was and b) to what extent the agent himself endorsed the killing himself. We show that people's willingness to blame the agent increases with the degree to which they judge that the agent could have done otherwise. Further, the more the agent endorses the killing, the more he is blamed. This effect is unaffected by the degree to which the agent lacks alternative possibilities. Following up, we explore which other situational constraints people accept to undermine free will and moral responsibility.

Is Lying Bound to What Is Said? Empirically Investigating Deceptive Presuppositions, Implicatures, and Actions.

Louisa Marie Reins and Alex Wiegmann

Lying is a moral phenomenon of considerable importance, as people are confronted with lies and engage in lying virtually every day. Many philosophical definitions of lying assume that for a person to lie, they need to explicitly state something they believe to be false. In this view, people cannot lie by merely presupposing or implicating believed-false content, nor by acting deceptively. Recent research, however, challenged this assumption by showing that false presuppositions and certain false implicatures are categorized as lies. As people do not seem to hold the definition outlined above, the present research aimed to identify the factors people instead take into consideration when categorizing different (speech) acts as lies, in order to get closer to an understanding of the folk concept of lying. To do so, we presented participants with descriptions of situations in which an agent deceived another person without explicitly stating something believed-false, each in four variants differing solely in how the misleading content was conveyed: as a presupposition, a generalized or particularized conversational implicature, or an action. After reading the stories, participants indicated the amount to which they thought each agent deceived and lied, and their amount of agreement to a set of six statements investigating factors hypothesized to be possible predictors of the lie ratings (order of the two question sets randomized between-subjects). The descriptive results from a pilot-study of the material ($N = 72$) suggest that all forms of misleading behaviour investigated were not only categorized as deceptions, but also clearly seen as lies. To our knowledge, our study is the first to show that people deem it possible to lie by acting deceptively. Furthermore, the lie judgments for all types of deceptions were highly correlated with how much the agent was thought to have committed themselves to the not explicitly-stated believed-false claim, by how contradictory it would be to try to cancel that claim, and by how hard it would be for the agent to talk themselves out of that claim ($r = .82 - .87$), with a combined measure of these three predictors correlating with the lie ratings even higher ($r = .89$). The calculability of the believed-false claim from the actual statement or action and the redundancy of reinforcing it, as well as whether the speaker wanted the claim to become common knowledge, were correlated with the lie ratings to a lesser extent ($r = .60 - .73$). Therefore, our preliminary results suggest that when judging whether different forms of deceptions are a lie, people seem to strongly consider whether the agent committed themselves to the believed-false claim, and therefore could not cancel that claim or talk themselves out of it easily. While this certainly is the case for prototypical instances of lying, we show that it can hold for deceptive presuppositions, implicatures, and even for deceptive actions as well. Final data collection and analysis will be completed by the time of the conference.

Looking Forward to Desert: Moral Responsibility and Forward-Looking Considerations

Jay Spitzley

Whether an agent is morally responsible for an action will depend on, among other things, his mental states at the time of the action. Insofar as being morally responsible for an action hinges on whether the agent deliberated about the action or was coerced into performing the action, facts pertaining to the time that this deliberation or coercion happened would also be considered relevant to moral responsibility. Part of deciphering which states and events are relevant to moral responsibility is determining the temporal relationship these relevant states and events have to an action. I call facts pertaining to states and events at or before the time of the action backward-looking considerations, and facts pertaining to states and events after the time of action forward-looking considerations. Although forward-looking considerations are commonly thought to be irrelevant to moral responsibility, I argue we have reason to question this presupposition. Contemporary philosophers agree that moral responsibility is fundamentally about desert. Although each account is different and some philosophers are more explicit than others about whether moral responsibility is entirely derivable from claims about desert, most take deserving or meriting praise and blame to play a central role in the nature of moral responsibility. It is also assumed that desert is not derived from consequentialist considerations, such as whether blaming an agent would deter them from performing similar bad acts in the future. What philosophers have not agreed on, historically speaking, is whether moral responsibility should be understood to be entirely backward-looking. In the past, many have argued that whether praise and blame are appropriate depends, at least in part, on forward-looking considerations, such as whether blame would produce certain beneficial consequences. Nonetheless, these consequentialist accounts of moral responsibility have fallen out of popularity because of their counterintuitive implications. Whereas the appeal of consequentialist accounts of moral responsibility comes from their future benefits (and not their intuitive nature), desert-based accounts rely on their intuitive appeal. To determine which specific desert-based account is correct, proponents of desert-based theories often employ the method of cases to draw conclusions regarding what is relevant to moral responsibility – a method by which the content of an intuition is treated as evidence either in favor of or against a proposed theory. In this paper, I provide empirical evidence that moral responsibility intuitions are sensitive to forward-looking considerations. This evidence not only poses a problem for views that rule out the relevance of forward-looking considerations, but also calls into question some popular presuppositions about moral responsibility. That is, if the method of cases is a permissible methodology for drawing conclusions about the nature of moral responsibility, then moral responsibility is not entirely backward-looking. Either philosophers have been wrong about the nature of desert and desert is not entirely backward-looking, or popular views of moral responsibility are mistaken and moral responsibility is not exclusively desert-based. This evidence also demonstrates intuitive support for consequentialist accounts of moral responsibility that have been largely abandoned due to their counterintuitive nature.

Folk Stories, Epistemic Normativity, and Intentional Action

Daniel Stermer

A much-debated area in contemporary philosophy of action is sorting out the nature of sideeffect actions. Here is one question that has been asked: which side-effect actions, if any, are intentional? There is a healthy amount of data for folk conceptions of “normal” intentional action. For example, the folk consider belief, desire, intention, skill, and awareness to be signs that an act is intentional (Malle and Knobe 1997). These properties nicely parallel conceptual analyses of intentional action argued for elsewhere (e.g., Mele and Moser 1994). In the case of side-effect actions, though, things are more difficult; it is unclear if and how each criterion applies to these special cases. Consider the oft-cited case of the chairman of a large corporation who undertakes a business venture which, as a side effect, harms the environment. Did the chairman intentionally harm the environment? Perhaps surprisingly, the folk answer “yes” (Knobe 2003). If the folk are on to something here, then this possibly undermines the intention and desire criteria cited above. There are at least four possible ways to explain the folk’s response: either the chairman’s indifference to harming the environment, the badness of the act itself, his confidence that starting the program would harm it, or some combination of these three. There is compelling data that the perceived rightness or wrongness of an action influences how the folk judge intentionality (Cushman and Mele 2007). In this paper, though, I want to focus on an underdeveloped area of experimental philosophy and intentional action: the role that epistemic normativity plays in side-effect actions— particularly, the influence that first-person determinations of side-effect action probability exert on folk judgments of intentionality. That is, there are cases in which we would expect an agent ought to know a side-effect action would occur in the course of performing some other action. Does such epistemic normativity have any bearing on the action’s intentionality? This loosely correlates to the aforementioned “confidence” criterion above but with the added factor of epistemic culpability. To answer this, I have created several vignettes and structured them as follows. First, I have separated the vignettes into two categories: those resulting in neutral side-effect actions and those resulting in bad side-effect actions. I have done this because, as the chairman case above indicates, moral considerations influence folk judgments of intentionality. Thus, including moral considerations helps to distinguish their effects from epistemic considerations. Furthermore, under both bad and neutral side-effect actions I have provided cases where the subjective probability (the agent’s judgment of probability) and objective (actual) probability match both high and low—that is, cases where the agent correctly judges the probability of some side-effect action occurring. Next, I provide cases where the agent misjudges the actual probability of a side effect action. The agent may do this in two ways: either she judges the objective probability to be high when it is low, or vice versa. This is where epistemic normativity comes into play. When the agent misjudges things, are they epistemically culpable? More importantly, was the side-effect action intentional? Based on the data, I conclude that epistemic considerations do influence folk to judge an action, when the agent is epistemically culpable, as intentional. I also offer philosophical considerations to support this conclusion.

When both can be right: Folk perceptions of moral disagreement and consequentialism on bioethical issues

Hugo Viciano

In this blitz talk and poster, I will first offer a quick glimpse into the results of an online nationally representative survey on attitudes of the adult Spanish population regarding a series of bioethical issues in the public agenda. One of the results of that survey was a robust association between so called folk moral relativistic attitudes and consequentialist attitudes towards the regulation of the bioethical issues.

That finding deserves further investigation and assessment, which is the main topic of this work. In particular, the operationalization of folk moral relativism has been criticized in a number of ways. To what extent does it represent genuine relativistic commitments? I bring evidence to disentangle the question of what is causing what in the association between relativistic attitudes and consequentialist attitudes. I will present the results of an attempt to measure the relative contributions of possible underlying factors such as social desirability, epistemic uncertainty, increased reflectiveness, and different dimensions of folk moral objectivism.

Philosophical Education Influences Responses to Thought Experiments: A Study of Two Non-Gettier Cases

Su Wu and Chengying Guan

Experimental philosophers find that our intuitions about philosophical thought experiments could vary among people with different backgrounds (Weinberg et al., 2001; Machery et al., 2004) and that intuitions are sensitive to unexpected factors such as the order of presentation of cases (Petrinovich & O'Neill, 1996; Swain et al., 2008) and participants' personality traits (Feltz & Cokely, 2009). On the assumption that intuitions have played a central role in philosophical study, experimental philosophers believe these findings constitute a serious challenge to the reliability of the armchair philosophical methodology.

However, whether we should accept the conclusion of these empirical studies depends on the actual role that intuitions play in armchair philosophy. Since armchair philosophers also appeal a lot to arguments in their work, it is controversial to assume that appealing to intuitions is all that armchair philosophy can do.

A typical discussion process in philosophy consists of both intuitions and arguments. What really matters is therefore not whether the intuitions themselves are unreliable. Instead, the crux is whether philosophers can adjust their unstable intuitions in the subsequent reflection processes, and finally come to establish tenable philosophical judgments.

It has been proposed that arguments may secure philosophers' responses to thought experiments (Deutsch, 2009; 2010; 2015; Cappelen, 2012), and intuitions should be modified by good arguments (Deutsch, 2010); this proposal needs to be verified by empirical studies. But some previous empirical results contradict this proposal, suggesting that arguments cannot modify responses to the Gettier case (Wysocki, 2016).

Nevertheless, there seem two defects in the previous study. For one thing, since the Gettier case is one of the very rare cases which philosophers shared an agreement while most thought experiments raise ambiguous intuitions (Cappelen, 2012), this consensus may lead to ceiling effect and undermine the influence of arguments. For the other, this study has not paid enough attention to "the expertise defense" which suggests a possibility of different responses' to arguments between people with and without philosophy background (Deutsch, 2010).

As a response to the problem just mentioned, we present a new empirical study based on two other cases in epistemology (the Truetemp case and the Barn case) which showed less consistency even among philosophers. In our study, 456 effective responses from Chinese participants were recycled online, including 254 participants who have taken at least one philosophy course. Replicating the existing view, our study reveals that arguments have no significant influence on participants' responses. However, on the other hand, we also find that arguments made a significant difference among subjects who have taken at least one philosophy course. Even participants with minor philosophical education can be more sensitive to arguments than participants without.

The conclusion of our study consists of three points. First of all, arguments work on participants who know philosophy, modifying their final judgments to issues considered in thought experiments. Secondly, it also reveals the limitation of experimental studies based only on questionnaires. The shortage of time and information in most previous experimental

philosophy studies hardly provide a condition for participants to fully understand the issue under discussion. Finally, this kind of full understanding may provide a new perspective to comprehend the improvement in philosophy and philosophical expertise.