INTUITIVE EXPERTISE IN MORAL JUDGEMENTS

Joachim Horvath* and Alex Wiegmann*

Ruhr University Bochum

* Authors contributed equally

This is a preprint of an article whose final and definitive form will be published in the Australasian Journal of Philosophy [2021]. The Australasian Journal of Philosophy is available online at: <http://www.tandf.co.uk/journals/>.

ABSTRACT

According to the ‘expertise defence’, experimental findings which suggest that intuitive judgements about hypothetical cases are influenced by philosophically irrelevant factors do not undermine their evidential use in (moral) philosophy. This defence assumes that philosophical experts are unlikely to be influenced by irrelevant factors. We discuss relevant findings from experimental metaphilosophy that largely tell against this assumption. To advance the debate, we present the most comprehensive experimental study of intuitive expertise in ethics to date, which tests five well-known biases of judgement and decision-making among expert ethicists and laypeople. We found that even expert ethicists are affected by some of these biases, but also that they enjoy a slight advantage over laypeople in some cases. We discuss the implications of these results for the expertise defence, and conclude that they still do not support the defence as it is typically presented in (moral) philosophy.

KEYWORDS: moral judgement; moral intuitions; expertise defence; framing effects; bias

1. Introduction

In this paper, we investigate whether professional moral philosophers are experts at making intuitive judgements about hypothetical moral cases. To advance the metaphilosophical discussion, we present new findings from a large and comprehensive online experiment performed with expert ethicists and lay subjects. In doing so, we will largely sidestep a controversial issue in the epistemology of ethics: the status of moral expertise and testimony. Ethicists tend to be sceptical of the possibility of gaining moral knowledge through (expert) testimony in the same way that we acquire knowledge from experts in, say, palaeontology or physics (see, for example, Hopkins [2007]; Lewis [2020]). However, certain limited forms of reliance on expert testimony seem possible in the moral domain as well (see, for example, Jones [1999]; Jones and Schroeter [2012]). We thus assume merely that professional ethicists may have intuitive expertise in their judgements about typical moral thought experiments. Paradigmatic examples would be Thomson's violinist [Thomson 1971] or the various trolley cases [Foot 1967; Thomson 1985]. Given that such thought experiments often involve unusual or even far-fetched scenarios, it seems *prima facie* plausible that the relevant intuitive judgements of professional ethicists are superior to those of laypeople. Therefore, intuitive moral expertise concerning ethical thought experiments should at least be a realistic possibility.

The question of whether expert (moral) philosophers really have such intuitive expertise has acquired special prominence in the metaphilosophical debate about experimental philosophy (see, for example, Weinberg et al. [2010]). Since findings from experimental philosophy suggest that intuitive judgements about hypothetical cases vary with philosophically irrelevant factors, some experimental philosophers have claimed that we should stop appealing to such judgements in philosophical theorizing (see, for example, Alexander et al. [2009]; Machery [2017]). Probably the most influential reply to this claim is

the *expertise defence* (see, for example, Horvath [2010]; Ludwig [2007]; Williamson [2011]), which is based on the observation that most studies from experimental philosophy are performed with philosophical laypeople. According to this defence, professional philosophers' intuitive judgements in their respective areas of expertise are superior to those of untrained laypeople, in analogy with other domains of academic expertise (compare Hales [2006]). The relevant claim for our purposes is that moral philosophers' intuitive judgements about ethical thought experiments can be expected to vary much less with morally irrelevant factors than those of laypeople. As a consequence, professional ethicists can ignore troubling experimental studies that are performed with lay subjects—and moral philosophy can carry on as before. Unfortunately, the empirical claim that moral philosophers enjoy intuitive expertise of the relevant kind has been undermined by recent studies in experimental metaphilosophy (see, for example, Wiegmann et al. [2020]). In the following, we will first present the empirical case against intuitive expertise in ethics, and then add our own, somewhat more complicated findings to the debate. Even in the light of our own findings, however, the prospects of the expertise defence in moral philosophy remain dim.

2. Intuitive Expertise and the Expertise Defence

Various findings from experimental philosophy suggest that intuitive judgements about hypothetical cases vary with philosophically irrelevant factors. For example, intuitive judgements about moral cases have been found to be sensitive to *order of presentation* (see, for example, Liao et al. [2012]; Wiegmann and Waldmann [2014]), even though the order of the cases has no obvious bearing on what is morally right or wrong in them (compare Wiegmann et al. [2020]). Other examples of arguably irrelevant factors are *cultural background* (see, for example, Machery et al. [2004]), *affective content* (see, for example,

Nichols and Knobe [2007]), and *heritable personality traits* (see, for example, Feltz and Cokely [2019]).

According to experimental philosophers of a *restrictionist* stripe, such findings suggest that intuitive judgements about philosophical thought experiments are untrustworthy (see, for example, Weinberg [2007]) or even flat-out unreliable (see, for example, Machery [2017]). If that is so, then the standard philosophical method of appealing to intuitions about hypothetical cases¹ should either be abandoned [Machery 2017], or at least seriously restricted until empirical research reveals which intuitions can be trusted [Weinberg 2015].

However, experimental restrictionism has met strong resistance, and the *expertise defence* has emerged as probably the most popular reply (see, for example, Alexander [2016]; Horvath [2010]; Ludwig [2007]; Machery [2017]; Nado [2014]; Weinberg et al. [2010]; Wiegmann et al. [2020]; Williamson [2011]). According to this defence, only the intuitive judgements of *philosophical experts*—that is, moral philosophers in the case of ethical thought experiments—matter for philosophical theorizing, just as in other respectable academic fields like mathematics or law (compare Hales [2006]). Since mathematicians typically ignore the mathematical intuitions of laypeople as irrelevant to their practice, professional philosophers can do likewise in their respective areas of expertise. This defence assumes that expert philosophers are less likely to be influenced by philosophically irrelevant factors than laypeople. If that is so, then expert philosophers can simply ignore problematic experimental results, which are usually obtained from philosophically untrained participants.

One important point is that the expertise defence is not concerned with the broader question of whether philosophers have some professional expertise or other, which they surely do—if only by knowing more about philosophy than laypeople. Rather, the key issue

¹ In recent metaphilosophy, there has been some controversy about whether it is in fact standard philosophical practice to appeal to intuitions about hypothetical cases (cf. Cappelen [2012]; Deutsch [2015]; Horvath [Unpublished Manuscript]; Horvath and Koch [2020]). For the purposes of this paper, we will ignore this strand of the metaphilosophical debate and simply assume that appealing to intuitions about cases as evidence is indeed standard philosophical practice—as continues to be the mainstream view in contemporary analytic philosophy.

is whether philosophers have specific intuitive expertise for judging hypothetical cases in their respective areas of philosophical competence, such as ethics or epistemology [Weinberg et al. 2010]. Even if there is a certain presumption that professional philosophers have this expertise [Horvath 2010; Williamson 2011], it is still an empirical question whether they actually have it—and experimental metaphilosophy has already begun to investigate this question. We will now briefly discuss the relevant findings, with a special focus on intuitive expertise in ethics.

3. Intuitive Expertise in (Moral) Philosophy

The seminal study of intuitive expertise in (moral) philosophy is by Schwitzgebel and Cushman [2012],² who found equally large *order effects* in expert ethicists and laypeople for standard ethical cases, including trolley cases of the type that we used in our own experiment (see below). Schwitzgebel and Cushman's study also had a high number of participants, and they used reasonably strict criteria for ethical expertise (= MA or PhD in philosophy, with ethics as an area of competence or specialization).

In a follow-up study, Schwitzgebel and Cushman [2015] replicated their earlier finding about order effects and found that professional philosophers were equally affected by Tversky and Kahneman's [1981] 'Asian disease' framing (see our own experiment below). Moreover, neither forced reflection, nor self-reported expertise, nor familiarity with the relevant cases made a difference to philosophers' susceptibility to the tested effects.

Tobia, Buckwalter, and Stich [2013] found that professional philosophers and lay people exhibit an equally large *actor-observer bias* concerning two well-known ethical cases, including a trolley case of the switch-type (see below). Actor-observer bias is the tendency to judge cases differently from a first-person and a third-person perspective. However, in a

² For a survey of empirical studies of intuitive expertise more generally, see Kahneman and Klein [2009].

follow-up study by Tobia, Chapman, and Stich [2013], there was no main effect of actor-observer bias. Instead, they found an interaction of this factor with a *cleanliness manipulation*, operationalized by the scent of Lysol, which is associated with cleanliness. Both of these effects should be taken with a grain of salt, though. First, although the actor-observer effects in Tobia, Buckwalter, and Stich [2013] were large, they were just below the conventional .05-threshold for statistical significance, and they also failed to replicate in Tobia, Chapman, and Stich [2013], as well as in our own experiment (see below). Second, the alleged surprising effects of subtle priming, as claimed for the scent of Lysol, have largely been discredited by the replication crisis of the last decade (see, for example, Open Science Collaboration [2015]).

A study by Löhr [2019] tested various hypotheses about Nozick's experience machine scenario with professional philosophers and laypeople. Even though professional philosophers were more consistent in their overall responses, 29% still gave inconsistent answers in a within-subject presentation of three variations of Nozick's scenario.

A recent study by Wiegmann, Horvath, and Meyer [2020] used a standard trolley dilemma and a six-option version of the same case, to test *order effects* and the influence of *irrelevant additional options* on moral judgements in laypeople and expert moral philosophers (= participants with a PhD or MA in philosophy, with ethics as an area of competence or specialization). Expert ethicists were no more resistant to order effects and irrelevant options than laypeople. Descriptively speaking, they were even more susceptible to these effects than laypeople (while also being more consistent in their responses).

Finally, let us briefly consider some findings on intuitive philosophical expertise outside of ethics. Hitchcock and Knobe [2009] found that the intuitions of philosophical experts about actual (or token) causation were affected by *norm violations* to the same extent

as those of laypeople. However, it is somewhat controversial whether this is really a philosophically irrelevant factor.

Schulz, Cokely, and Feltz [2011] found that the facet *warmth* of the personality trait *extraversion* predicts compatibilist intuitions about free will in both philosophical experts and laypeople (see also Feltz and Cokely [2019]).

In a study performed with language experts from subfields of philosophy and linguistics, Machery [2012] found that experts have no clear advantage over laypeople. Moreover, the various expert subgroups *disagreed with each other* in their intuitive judgements about reference.

Horvath and Wiegmann [2016] compared expert epistemologists' and laypeople's intuitive judgements about knowledge cases and found significant expert-lay differences in a number of cases. However, these differences were all quantitative rather than qualitative (that is, laypeople agreed significantly more strongly with the tested knowledge-claims than experts, but the general tendency was the same). More strikingly, expert epistemologists sometimes disagreed with the 'textbook consensus' in epistemology, for example, by judging fake-barn-type cases to be instances of knowledge (as do laypeople).

In a recent study, Schindler and Saint-Germier [Unpublished Manuscript] compared intuitive judgements and the relevant interpretative abilities of professional philosophers and laypeople concerning standard thought experiments from theoretical philosophy. Overall, they found that philosophers have an advantage over laypeople vis-à-vis the relevant 'textbook consensus' in only three out of six cases.³

The upshot of the reviewed studies is that philosophical experts at best enjoy a mild advantage over laypeople in a few of the tested cases. The empirical data thus do not support the expertise defence as a convincing reply to the experimental restrictionist challenge.

³ Further studies that bear on the expertise defence less directly are Beebe and Monaghan [2018], Carter et al. [2016], Carter et al. [2019], and Sytsma and Machery [2010].

Moreover, the prospect of a general, unqualified expertise defence ‘from the philosophical armchair’ seems near hopeless at this point, because the cases where philosophical experts may have a genuine advantage over laypeople are few and far between, and cannot be predicted in advance of empirical investigation.

However, it would be premature to conclude that the expertise defence has been empirically refuted. First, the number of investigated biases in the moral domain (= 5) is still fairly small.⁴ Second, given that professional philosophers are difficult to recruit and surprising findings are much easier to publish, researchers may have primarily studied biases that they expected ‘to work’ with expert participants. The current literature may therefore paint an overly negative picture of (moral) philosophers’ relevant expertise. To broaden and ‘debias’ this research, our own study investigates as many biases as all the previous studies combined (= 5), which we have, for the most part, chosen independently of specific expectations concerning expert participants (see below for details). We will now present the results of this large and comprehensive online experiment performed with expert moral philosophers and laypeople.

4. Experiment

The preregistration of our experiment can be found here:

https://osf.io/t5kpa/?view_only=aca56e761a9e4e6f9b9ee7f0197594d4

4.1 Participants

Our analysis includes the data of 454 participants. Their mean age was 36 years. 53% identified as *male* and 45% as *female*, while 2% selected *other* or *prefer not to say*. 227 were

⁴ These are *order effects* [Schwitzgebel and Cushman 2012, 2015; Wiegmann et al. 2020], ‘*Asian disease*’ *framing* [Schwitzgebel and Cushman 2015], *actor-observer bias* [Löhr 2019; Tobia, Buckwalter, et al. 2013; Tobia, Chapman, et al. 2013], *cleanliness priming* [Tobia, Chapman, et al. 2013], and *irrelevant additional options* [Wiegmann et al. 2020].

expert participants (mean age 37 years, 76% *male*, 21% *female*, and about 3% *other* or *prefer not to say*), and equally many, 227, were lay participants (mean age 35 years, 39% *male*, 69% *female*, and 1% *other* or *prefer not to say*).

Expert participants were recruited through a call for participation with a link to the online experiment, which was distributed on an electronic mailing list for philosophers (PHILOS-L: <https://www.liverpool.ac.uk/philosophy/philos-l>), a blog specialized in ethics (PEA Soup: <http://peasoup.us>), and two social media platforms (Facebook and Twitter). To qualify as experts, participants had to indicate having a PhD (59%) or an MA (41%) in philosophy, with moral philosophy or ethics as one of their areas of specialization (51%) or competence (49%). We excluded the data of participants who did not complete the survey, or took less than 2 minutes to do so, or failed to answer both attention checks correctly. 227 out of 1033 participants met these criteria, which were not mentioned in the call for participation or during the survey.⁵

Lay participants were recruited on Prolific Academics, a database based in the UK (compare Palan and Schitter [2018]). Each participant received compensation of £0.40. To ensure that lay participants differed strongly from expert participants in their level of philosophical expertise, we only included the data of participants who had either no prior experience with philosophy, or only some experience in school or university (without a degree in philosophy), and who were not current students of philosophy. We also excluded the data of participants who did not complete the survey, or took less than 2 minutes to do so, or failed to answer both attention checks correctly.⁶ 227 out of 336 participants met these criteria, which were not mentioned in the call for participation or during the survey.⁷

⁵ The main reasons for this high exclusion rate were not having a PhD or MA in philosophy (639 participants), and not having moral philosophy/ethics as an area of specialization or competence (563 participants).

⁶ 12% of expert and 19% of lay participants failed to answer both attention checks correctly.

⁷ In line with our preregistration, lay participants were recruited until we reached the number of valid expert participants at the end of May 2019.

4.2 Design and Materials

The experiment was conducted online on Unipark (<https://www.unipark.com/en/survey-software/>). Participants were first presented with general instructions that familiarized them with the question mode, asked them to read the case descriptions carefully, and appealed to them to take the task seriously. Participants were then randomly assigned to one of two *framing direction* conditions (Positive or Negative), and presented with five moral scenarios that we used to implement five different *framing effects* (Focus, Prospect, Accounting, Perspective, and Decoy). The five scenarios were framed in a way expected to either increase (in the Positive condition) or decrease agreement with the presented moral claims (in the Negative condition). This resulted in a 2*2*5 mixed design, with *framing direction* (Positive vs. Negative) and *level of expertise* (Expert vs. Lay) as between-subjects factors, and *framing effect* (Focus, Prospect, Accounting, Perspective, and Decoy) as within-subject factor. The order of presentation was fixed, as presented below.⁸ Two attention checks concerning the Accounting and Decoy scenario were included (see below).

A pilot study with lay participants suggested that the *framing direction* manipulation works in the expected direction, with effects ranging from small to large across vignettes. One might worry that small or even non-significant effects in laypeople are not useful for testing the expertise defence, because experts can hardly outperform laypeople here by being *less biased* than the latter. This worry is misplaced, however. First, pilot studies provide only rough estimates of effect sizes, and thus predict only roughly what will happen in the main study. Second, there is still room for experts to outperform laypeople in individual scenarios, and also averaged across all scenarios. Moreover, experts are sometimes even more biased

⁸ This order of presentation was chosen to minimize potential order effects. For instance, it has been shown that presenting a ‘push-type’ dilemma can strongly influence the evaluation of subsequent scenarios. This is why we presented Decoy, which includes a push-type option, as the final scenario. Moreover, the two structurally similar scenarios in which a threat can be redirected by the agent (Focus and Perspective) were not presented one right after the other, but rather with two scenarios in between (Prospect and Accounting).

than laypeople (see, for example, Chi [2006]; Wiegmann et al. [2020]). Third, and most importantly, selecting only scenarios that yield strong biases in laypeople stacks the deck in favour of the experts, for it makes it relatively easy for them to outperform laypeople in the experimental setting, even though they might not really be better on average. To illustrate, consider a teacher who wants to compare the performance of two students A and B, and does so by first testing A on a set of questions, and then asking B only those questions that A got wrong. This flawed procedure would make it very easy for B to outperform A, because B only needs to get a few of the questions right that A got wrong, while at the same time it is completely screened off how B would perform on the questions that A got right. Therefore, we did not dismiss any scenario simply because the pilot study indicated only a small effect, and we selected most of the original scenarios independently of expectations concerning philosophical experts (with the exception of Prospect and Perspective; see below).

Here is the exact wording of our five scenarios, with the relevant manipulations in square brackets (with bold font indicating the Positive condition, for which higher ratings were expected):

Focus (question-focus on good vs. bad consequences; Petrinovich and O'Neill [1996])

Carl is standing on the pier and observes a shark swimming toward five swimmers. There is only one possibility to avoid the death of the five swimmers: Carl could make loud noises, and thereby divert the shark into his direction. However, there is a fisherman swimming on the path between Carl and the shark. Diverting the shark would therefore save the five swimmers but kill the fisherman.

How much do you disagree or agree with the following claim:

Carl should make loud noises, which will result in [**the five swimmers being saved** / the fisherman being killed].

This manipulation of the question-focus in terms of good or bad consequences is based on Petrinovich and O'Neill [1996], who originally used the trolley paradigm to demonstrate its effect. In our adapted version, the trolley/train is replaced by a shark in order to avoid interferences with our second trolley-style scenario (see below). Originally, this manipulation resulted in a large ($\eta_p^2 = .45$) and robust effect (39 out of 40 paired cases received significantly lower agreement ratings for the kill-wording). Moreover, the applied framing is clearly morally irrelevant, for the potential action and its consequences are exactly identical.

Prospect (prospect framed in positive vs. negative terms; Tversky and Kahneman [1981])

A country is preparing for the outbreak of an unusual disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows:

If Program A is adopted, [**200 people will be saved** / 400 people will die].

If Program B is adopted, there is a one-third probability that [**600 people will be saved** / nobody will die] and a two-thirds probability that [**no people will be saved** / 600 people will die].

How much do you disagree or agree with the following claim:

Program A should be adopted.

This vignette tests a seminal effect known as the ‘Asian Disease Problem’, introduced by Tversky and Kahneman [1981], which is one of a variety of effects that show that identical prospects are evaluated differently when framed in terms of losses versus gains. The original scenario is very similar to the one we used—only the original test question (‘Which of the two programs would you prefer?’) was modified to highlight the moral dimension of the scenario. The effect found in the original study was large—Program A was favoured by 72% of the participants when framed in terms of saving, as compared to 22% when framed in terms of dying—and it was also found to be highly robust (see, for example, Ruggeri et al. [2019], who replicated effects of this kind in nineteen countries). Since the two framings are identical in all morally relevant aspects, an effect of this kind would clearly indicate biased moral judgements. In fact, the ‘Asian Disease Problem’ has already been tested among professional philosophers (with only slightly different vignettes), and they were not found to be immune to this effect (Schwitzgebel and Cushman [2015]). A replication with our ethical version would thus provide converging evidence that even professional philosophers (in our case: moral philosophers) are susceptible to this seminal effect.

Accounting (mental accounting; Thaler [1985]; Tversky and Kahneman [1981])

Sarah has decided to go to a charity event [**where admission will be \$20 per ticket** / and already bought a ticket for \$20]. All of those \$20 will be spent on improving the living conditions in some of the world’s poorest countries, but only if the ticket is signed by the ticket holder and handed over to a representative of

the charity organization during the event. As Sarah enters the venue where the charity event takes place, she discovers that she has [lost a \$20 note / lost her ticket] on the way to the venue. Frustrated with her loss, Sarah decides against spending \$20 on [a / another] ticket.

How much do you disagree or agree with the following claim:

Sarah should be blamed for not spending \$20 on [a / another] ticket.

[Attention check, on next page:]

What did Sarah lose on her way to the venue? Her ticket/Her mobile phone/A \$20 note

This vignette is a moral adaptation of another seminal finding by Tversky and Kahneman [1981], and Thaler [1985]. The key idea is that we have different mental accounts for different kinds of costs, which can lead to different decisions in financially identical situations. In the original version, participants are asked to imagine that they had decided to see a play whose admission is \$10 per ticket. In one condition, they were told that they had bought a ticket only to discover that they had lost it. In another condition, they had not bought a ticket, but instead lost a \$10 bill on their way to the theatre. When asked whether they would still pay \$10 for a ticket, almost all participants (88%, Tversky and Kahneman [1981]) answered ‘Yes’ in the condition where they lost a \$10 bill, in contrast to only about half of them (46%) in the condition where they lost the original ticket. We adapted this paradigm to the moral domain by stipulating that the money for the ticket goes to a charity devoted to fighting extreme

poverty. The original findings suggest lower blame ratings in the condition where the ticket was lost—under the assumption that people tend to blame other people less if they themselves would have behaved similarly. Without making any questionable additional assumptions—for example, that losing a ticket is more frustrating than losing a \$10 bill and that more frustrated people deserve less blame—it is hard to see how these two versions could differ in morally relevant aspects.

Perspective (actor vs. observer perspective; Tobia, Buckwalter, et al. [2013]; Tobia, Chapman, et al. [2013])

[**Anne is** / You are] a technician in a company producing chemicals and responsible for the safety of employees. One day, [**she observes** / you observe] on [**her** / your] monitor that poisonous gas has been released into the ventilation system and is about to contaminate a room with three researchers, who will die if they are exposed to the gas. Because there is not enough time to warn the three researchers, there is only one possibility to avoid their death: [**Anne is** / You are] in control of the ventilation system and could close the ventilation shaft leading to the room with the three researchers, who would be saved by this measure.

However, as a result of closing the ventilation shaft, the poisonous gas would be released into another room where one researcher is working. Since there is no way to block the ventilation shaft to this other room and also not enough time for a warning, the one researcher working there would die from being exposed to the gas.

How much do you disagree or agree with the following claim:

[**Anne** / You] should close the ventilation shaft.

This vignette tests the finding by Tobia, Buckwalter, and Stich [2013] that professional philosophers are equally susceptible to actor-observer bias in moral scenarios as laypeople. Actor-observer bias trades on the contrast between judging decisions from the first-person perspective of the person who makes the decision (the actor-perspective), and a more neutral third-person perspective (the observer-perspective). One reason for including this alleged bias was to clarify earlier findings by Tobia and colleagues [2013, 2013], which are somewhat mixed (see above). One puzzling aspect of their results is that the actor-observer manipulation affected professional philosophers and laypeople in opposite directions. For instance, in a switch-type trolley dilemma (Tobia, Buckwalter, et al. [2013]; experiment 2), the observer framing increased permissibility ratings in laypeople (with respect to redirecting the trolley), but it decreased permissibility ratings in professional philosophers. Since we are consistently re-phrasing the test question in terms of whether a certain decision *should* be made (in contrast to previous studies), it was not antecedently clear how our manipulation would influence participants' ratings. Thus, we conducted a pilot study that indicated that laypeople's agreement ratings should be higher in the observer condition. We also made the conservative assumption that the direction of the effect, if any, would be the same for experts and laypeople. An actor-observer effect would clearly indicate a bias in our scenario, because it seems *ceteris paribus* morally irrelevant who makes the decision here.

Decoy (decoy effect; Huber et al. [1982])

On the test ground of a modern railroad property, an unmanned speed-train that normally can be remote-controlled is out of control due to a technical defect.

This speed-train is heading towards five railroad workers that are maintaining the tracks (see figure below). Since these workers are wearing a new type of hearing protection, they would not notice the speed-train on time, and hence would be run over by it and lose their lives in the accident. Peter, an employee of the rail track control center, recognizes the upcoming accident.

[Positive condition (three options):]

Peter can choose between exactly three options:

Option 1: Peter could do nothing. The five workers would be run over by the speed-train and lose their lives in the accident.

Option 2: Peter could push a button that would open a trap door installed in the right bridge and cause the two workers on top of the right bridge to fall on the tracks. The speed-train would collide with the two workers and be stopped before it reaches the five workers. The two workers would lose their lives due to the collision.

Option 3: Peter could push a button that would open a trap door installed in the left bridge and cause the one worker on top of the left bridge to fall on the tracks. The speed-train would collide with the one worker and be stopped before it reaches the five workers. The one worker would lose his life due to the collision.

How much do you disagree or agree with the following claim:

Peter should choose Option 3, that is, he should push the button that will open the trap door in the left bridge and cause the one worker to fall on the tracks.

[Negative condition (two options):]

Peter can choose between exactly two options:

Option 1: Peter could do nothing. The five workers would be run over by the speed-train and lose their lives in the accident.

Option 2: Peter could push a button that would open a trap door installed in the left bridge and cause the one worker on top of the left bridge to fall on the tracks.

The speed-train would collide with the one worker and be stopped before it reaches the five workers. The one worker would lose his life due to the collision.

How much do you disagree or agree with the following claim:

Peter should choose Option 2, that is, he should push the button that will open the trap door in the left bridge and cause the one worker to fall on the tracks.

[Attention check (next page):]

How many bridges were displayed in the figure on the previous page? 1/2/3

This vignette tests the ‘decoy effect’, one of the best-known human violations of rational choice (also known as the ‘attraction effect’ or ‘asymmetric dominance effect’; Huber et al. [1982]; Huber and Puto [1983]; see Huber et al. [2014], for an overview; see Frederick et al. [2014], for boundary conditions). This effect consists of adding an inferior option to a choice-set which increases the choice-share of the option that it most closely resembles. In our moral adaptation, the additional option that results in the death of two people should thus increase the ratings for the otherwise identical option where only one person dies.

The response format we used was the same for all five scenarios: a 7-point Likert-item ranging from ‘strongly disagree’ (1) to ‘strongly agree’ (7).

4.3 Results

Data available online at:

https://osf.io/y5x6b/?view_only=1d233da8938c414da95cb1288884ea76

Our results are summarized in *Figure 1*, *Table 1*, and *Table 2*. We will first analyse the overall effect of *framing direction* and its interaction with *level of expertise*, and then consider each *framing effect* separately.

Figure 1. Participants' agreement ratings as a function of *framing effect*, *level of expertise*, and *framing direction*. Error bars represent 95% confidence intervals. Response options range from 1 ('strongly disagree') to 7 ('strongly agree').

4.3.1 Global Analyses

We applied a mixed ANOVA with *framing direction* (Positive vs. Negative) and *level of expertise* (Expert vs. Lay) as between-subjects factors, and *framing effect* as within-subject factor. All results are summarized in *Table 1*; we will now discuss the most relevant ones.

The *framing direction* manipulation (Positive vs. Negative) affected participants' responses in the expected direction, with mean ratings in the Positive condition ($M=4.49$, $SE=0.07$) being significantly higher than in the Negative condition ($M=4.00$, $SE=0.07$), $F(1, 450) = 25.68$, $p < .001$, $\eta_p^2 = .054$. The interaction of *level of expertise* and *framing direction*

was not significant, $F(1, 450) = 2.56$, $p = .111$, $\eta_p^2 = .0056$, indicating that the overall effect of *framing direction* did not differ significantly between moral philosophers and laypeople.⁹

The interaction of *framing effect* and *framing direction* was significant, indicating that the five framing effects differed in how strongly they affected participants' responses, $F(4, 1800) = 6.02$, $p < .001$, $\eta_p^2 = .013$.

Lastly, the interaction of *framing effect*, *level of expertise*, and *framing direction* was not significant (albeit almost), $F(4, 1800) = .22$, $p = .064$, $\eta_p^2 = .005$, indicating that the pattern of effects caused by the *framing direction* manipulation in the five scenarios was roughly the same for lay and expert participants. However, this still leaves open whether individual framing effects affected moral philosophers and laypeople differently (see below for further analysis).

Table 1. Results of mixed ANOVA with *level of expertise* (Expert vs. Lay) and *framing direction* (Positive vs. Negative) as between-subjects factors, and *framing effect* (Focus, Prospect, Accounting, Perspective, Decoy) as within-subject factor.

Effect	DF	F	p	η_p^2
Level of expertise	1	0.31	.577	.001
Framing direction	1	25.68	<.001	.054
Level of expertise * Framing direction	1	2.56	.111	.006
Error	450			
Framing effect	4	91.86	<.000	.170
Framing * Level of expertise	4	0.51	.731	.001
Framing effect * Level of expertise	4	6.02	.000	.013
Framing effect * Level of expertise * Framing direction	4	2.22	.064	.005
Error	1800			

Limiting the analysis to moral philosophers (Expert), a repeated-measure ANOVA with *framing direction* (Positive vs. Negative) as between-subjects factor and *framing effect* as

⁹ A linear mixed model analysis with *participant* and *framing effect* as random effect led to the same results—as expected given the specifics of our design (balanced, no missing values, etc.). This equivalence also holds for the following analyses.

within-subject factor indicated that, across all scenarios, the *framing direction* manipulation affected the agreement ratings of experts in the expected direction. Mean ratings in the Positive condition (M=4.38, SE=0.11) were significantly higher (on average by 0.33 points) than in the Negative condition (M=4.05, SE=0.11), $F(1, 225) = 4.95$, $p < .027$, $\eta_p^2 = .0215$.

Limiting the analysis to laypeople (Lay), the analogous ANOVA led to similar results. Mean ratings in the Positive condition (M=4.59, SE=0.09) were significantly higher (on average by 0.64 points) than in the Negative condition (M=3.95, SE=0.08), $F(1, 225) = 28.34$, $p < .001$, $\eta_p^2 = .112$.

To conclude, our *framing direction* manipulation affected both expert moral philosophers and laypeople equally, with no significant difference between these two groups.¹⁰ Descriptively, however, the effect size of *framing direction* was only small-to-medium for experts, but medium-to-large for laypeople.¹¹

4.3.2 Analyses at the Level of Individual Framing Effects

Table 2 displays the effect of *framing direction* (Positive vs. Negative) and *level of expertise* (Expert vs. Lay) for each of the five *framing effects*. We chose (and preregistered: https://osf.io/t5kpa/?view_only=aca56e761a9e4e6f9b9ee7f0197594d4) this as our primary level of analysis for the following reason: the five framings influence participants' responses differently, even if some of them look superficially similar. For instance, both question-focus and prospect-framing manipulate whether the word 'killing' or 'saving' is used, but the underlying psychological mechanism is still different. The manipulation of question-focus merely highlights the negative/positive aspect of the available action, while prospect-framing makes people prefer risky/safe choices if a prospect is framed negatively/positively.

¹⁰ Given that the gender distribution differed strongly between expert and lay participants, we also tested the interaction of *gender* and *framing direction*, which was not significant, $p = .76$ (i.e. the *framing direction* affected both genders equally).

¹¹ The conventions for reported effect sizes are as follows: 0.01 = small, 0.06 = medium, 0.14 = large.

Therefore, the influence of the five effects might differ as well,¹² and we wanted to find out how each individual framing affects our moral intuitions.¹³ Some of the framings might also influence expert and lay participants differently, which can only be discovered by analysing each framing condition separately.¹⁴

Given our preregistered significance criterion ($p < .05$, two-sided), lay responses were significantly influenced by two of the five framing effects (Focus and Prospect), while expert responses were only significantly affected by one framing effect (Prospect). Applying a less strict significance criterion ($p < .05$, one-sided) would render two additional framing effects (Accounting and Decoy) significant for both moral philosophers and laypeople, and so experts would be affected by morally irrelevant factors in three out of five cases, and laypeople in four out of five cases.

Concerning expert-lay differences at the level of individual framing effects, moral philosophers were thus significantly less influenced by two framing effects (Focus and Prospect), while there was no significant expert-lay difference concerning the other three effects.

¹² This expectation was confirmed by the significant interaction of *framing direction* and *framing effect*.

¹³ The distinctness of the five framing effects is also why p-values in *Table 2* are not adjusted for multiple comparisons.

¹⁴ This possibility obtains irrespective of the fact that the overall interaction of *framing direction*, *framing effect*, and *level of expertise* was not significant ($p = .064$).

Table 2. Number of participants, means, 95% confidence interval, and standard deviation for all combinations of *framing direction*, *framing effect* and *level of expertise*. The columns for p-values and effect sizes (one-way ANOVA) concern the *framing direction* manipulation (Positive vs. Negative) in the respective group of moral philosophers (Expert) or laypeople (Lay), for each of the five framing effects (Focus, Prospect, Accounting, Perspective, Decoy). P-values and effect sizes in the ‘Expert vs. Lay’ column concern the comparison of moral philosophers and laypeople in relation to the *framing direction* manipulation in the respective framing condition, that is, the interaction of *level of expertise* and *framing direction* for the *framing effect* at issue. Reported p-values are not adjusted (see footnote 13).

Framing Effect	Expertise	Framing Direction	N	Mean	95% CI	Std	p-value	Effect (η²)	Expert vs. Lay (p and η² _p)	
Focus	Expert	Positive	113	4,84	[4,51; 5,17]	1,77	.506	.002	.018*	.012
Focus	Expert	Negative	114	4,69	[4,4; 4,98]	1,57				
Focus	Lay	Positive	113	5,12	[4,89; 5,36]	1,25	<.001*	.081		
Focus	Lay	Negative	114	4,29	[4; 4,58]	1,55				
Prospect	Expert	Positive	113	4,19	[3,86; 4,51]	1,75	.007*	.032	.006*	.017
Prospect	Expert	Negative	114	3,59	[3,3; 3,88]	1,55				
Prospect	Lay	Positive	113	4,61	[4,29; 4,93]	1,73	<001*	.173		
Prospect	Lay	Negative	114	3,18	[2,91; 3,44]	1,41				
Accounting	Expert	Positive	113	3,48	[3,11; 3,84]	1,96	.058	.016	.981	~0
Accounting	Expert	Negative	114	3,00	[2,66; 3,34]	1,82				
Accounting	Lay	Positive	113	3,49	[3,1; 3,87]	2,07	.077	.014		
Accounting	Lay	Negative	114	3,02	[2,66; 3,37]	1,90				
Perspective	Expert	Positive	113	4,84	[4,51; 5,17]	1,78	.910	~0	.851	~0
Perspective	Expert	Negative	114	4,82	[4,53; 5,1]	1,54				
Perspective	Lay	Positive	113	4,98	[4,72; 5,24]	1,41	.668	~0		
Perspective	Lay	Negative	114	4,90	[4,65; 5,16]	1,36				
Decoy	Expert	Positive	113	4,57	[4,2; 4,93]	1,98	.092	.013	.920	~0
Decoy	Expert	Negative	114	4,15	[3,83; 4,47]	1,73				
Decoy	Lay	Positive	113	4,74	[4,43; 5,06]	1,70	.087	.013		
Decoy	Lay	Negative	114	4,36	[4,05; 4,67]	1,66				

5. Discussion

We tested five well-known cognitive biases in our experiment, such as ‘Asian disease’ framing and the decoy effect, and we also placed special emphasis on their replicability. The only exception is actor-observer framing, which we included on the basis of previous results with professional philosophers [Tobia, Buckwalter, et al. 2013; Tobia, Chapman, et al. 2013]. Our own results are based on a much higher number of participants, and they strongly suggest that actor-observer framing in moral scenarios is not a robust effect (and is perhaps even a non-effect) for both philosophical experts and laypeople, which is in line with another failed replication of actor-observer framing (reported in Cova et al. [2018]).

Since many of the tested biases are also well-known in ethics, it is remarkable that moral philosophers were still significantly affected by these biases taken together, that is, over all five scenarios. Even at the level of individual scenarios, moral philosophers were significantly affected by Tversky and Kahneman’s ‘Asian disease’ framing, which is one of the most famous psychological findings. Therefore, bias was neither eliminated by expertise in ethics, nor by the widespread familiarity of the tested effects. Adding to many previous results (see section 3), our findings thus cast further doubt on the expertise defence against the challenge from experimental restrictionism.

However, our results are also more mixed and complicated than those of previous studies. For example, expert moral philosophers—unlike laypeople—were not significantly affected by a simple framing in terms of people *killed* versus people *saved* in our Focus scenario. And while expert ethicists were significantly influenced by our ‘Asian disease’ framing in the Prospect case, they were still significantly less affected than laypeople (see above). Our results therefore indicate that expert ethicists have a genuine advantage over laypeople with respect to some well-known biases.

So, is our study good news for proponents of the expertise defence? That depends on what their ambition is. If they aim for a general ‘armchair’ defence of appeals to intuitions about cases in ethics, then our mixed findings are not exactly good news for them. For, whether moral philosophers are more immune to bias in a given case can neither be settled from one’s philosophical armchair, nor by some very general argument in favour of moral philosophers’ expertise. Instead, we can only tell whether moral philosophers are less affected by certain biases once we have carried out sufficient experimental research. Very general armchair versions of the expertise defence are thus undermined by our mixed results to the same extent that they were by the more clear-cut previous findings (compare section 3).

However, if (moral) philosophers are open to pursuing a more empirical and piecemeal version of the expertise defence, then our findings may provide some reasons for optimism. For example, the fact that expert ethicists don’t fall for a simple saving/killing framing in our Focus case might be a reason to trust expert intuitions more when it comes to trolley scenarios of the switch-variety. One key issue for such a piecemeal empirical approach to the expertise defence would be whether the relevant findings can be generalized to other (similar) cases or (related) biases. Absent further experimental data, it might be unclear, for example, whether moral philosophers are equally immune to a saving/killing framing in trolley scenarios of the push-type. Therefore, the ambitions of an empirical expertise defence may have to be very modest indeed, and thus are bound to disappoint most actual proponents of the expertise defence in (moral) philosophy.

6. Conclusion

We first considered the experimental restrictionist challenge to intuitions about cases, with a special focus on moral philosophy, and then introduced the expertise defence as the most popular reply. The expertise defence makes the empirically testable assumption that the case intuitions of expert philosophers are significantly less influenced by philosophically irrelevant factors than those of laypeople. The upshot of our discussion of relevant findings from experimental metaphilosophy was twofold: first, extant findings largely tell against the expertise defence, and second, the number of published studies and investigated biases is still fairly small. To advance the debate about the expertise defence in moral philosophy, we thus tested five well-known biases of judgement and decision-making among expert ethicists and laypeople. Averaged across all biases and scenarios, the intuitive judgements of both experts and laypeople were clearly susceptible to bias. However, moral philosophers were also less biased in two of the five cases (Focus and Prospect), although we found no significant expert-lay differences in the remaining three cases.

In comparison to previous findings (for example Schwitzgebel and Cushman [2012, 2015]; Wiegmann et al. [2020]), our results appear to be relatively good news for the expertise defence, because they suggest that moral philosophers are less influenced by some morally irrelevant factors, such as a simple saving/killing framing. On the other hand, our study does not support the very general armchair versions of the expertise defence that one often finds in metaphilosophy, which try to reassure (moral) philosophers that they need not worry about the influence of philosophically irrelevant factors. At best, however, we need not worry about just a few cases and a few human biases—and even that modest hypothesis can only be upheld on the basis of sufficient empirical research.¹⁵

¹⁵ We would like to thank John Horden, Steffen Koch, Michael Vollmer, and several anonymous reviewers for very helpful comments on previous versions of this paper. We would also like to thank our audiences at *EuroCogSci*, Ruhr University Bochum, September 2019, at the *Experimental Philosophy Conference Bern*,

Funding

Joachim Horvath and Alex Wiegmann's work on this paper was supported by an Emmy Noether grant of the *German Research Foundation* (DFG), project number 391304769.

References

- Alexander, J. 2016. Philosophical Expertise, in *A Companion to Experimental Philosophy*, eds. J. Sytsma and W. Buckwalter, Malden, MA: Wiley Blackwell: 555–67.
- Alexander, J., R. Mallon, and J.M. Weinberg 2009. Accentuate the Negative, *Review of Philosophy and Psychology*, 1/2: 297–314.
- Beebe, J.R., and J. Monaghan 2018. Epistemic Closure in Folk Epistemology, in *Oxford Studies in Experimental Philosophy, Volume Two*, eds. J. Knobe, T. Lombrozo, and S. Nichols, Oxford: Oxford University Press: 38–70.
- Cappelen, H. 2012. *Philosophy Without Intuitions*, Oxford: Oxford University Press.
- Carter, J.A., M. Peterson, and B. van Bezooijen 2016. Not Knowing a Cat is a Cat: Analyticity and Knowledge Ascriptions, *Review of Philosophy and Psychology*, 7/2: 817–34.
- Carter, J.A., D. Pritchard, and J. Shepherd 2019. Knowledge-how, Understanding-why and Epistemic Luck: An Experimental Study, *Review of Philosophy and Psychology*, 10/4: 701–34.
- Chi, M. 2006. Two Approaches to the Study of Experts' Characteristics, in *The Cambridge Handbook of Expertise and Expert Performance*, eds. K.A. Ericsson, N. Charness, P. Feltovich, and R. Hoffman, Cambridge: Cambridge University Press: 21–30.

University of Bern, September 2019, and at the *Tagung für Praktische Philosophie*, Panel *Gedankenexperimente in der praktischen Philosophie*, University of Salzburg, October 2020. Last but not least, many thanks to Nick Byrd, Daniel Dennett, and Peter Singer for distributing our call for participation on Twitter.

- Cova, F., B. Strickland, A. Abatista, et al. 2018. Estimating the Reproducibility of Experimental Philosophy. *Review of Philosophy and Psychology*, URL = <<https://doi.org/10.1007/s13164-018-0400-9>>
- Deutsch, M. 2015. *The Myth of the Intuitive: Experimental Philosophy and Philosophical Method*, Cambridge, MA: The MIT Press.
- Feltz, A., and E. Cokely 2019. Extraversion and Compatibilist Intuitions: A Ten-year Retrospective and Meta-analyses, *Philosophical Psychology*, 32/3: 388–403.
- Foot, P. 1967. The Problem of Abortion and the Doctrine of Double Effect, *Oxford Review*, 5: 5–15.
- Frederick, S., L. Lee, and E. Baskin 2014. The Limits of Attraction, *Journal of Marketing Research*, 51/4: 487–507.
- Hales, S. 2006. *Relativism and the Foundations of Philosophy*, Cambridge, MA: The MIT Press.
- Hitchcock, C., and J. Knobe 2009. Cause and Norm, *The Journal of Philosophy*, 106/11: 587–612.
- Hopkins, R. 2007. What Is Wrong with Moral Testimony?, *Philosophy and Phenomenological Research*, 74/3: 611–34.
- Horvath, J. 2010. How (Not) to React to Experimental Philosophy, *Philosophical Psychology*, 23/4: 447–80.
- Horvath, J. Unpublished Manuscript. *Mischaracterization Reconsidered*.
- Horvath, J., and S. Koch 2020. Experimental Philosophy and the Method of Cases, *Philosophy Compass*, URL = <<https://doi.org/10.1111/phc3.12716>>
- Horvath, J., and A. Wiegmann 2016. Intuitive Expertise and Intuitions about Knowledge, *Philosophical Studies*, 173/10: 2701–26.
- Huber, J., J.W. Payne, and C. Puto 1982. Adding Asymmetrically Dominated Alternatives:

- Violations of Regularity and the Similarity Hypothesis, *Journal of Consumer Research*, 9/1: 90–8.
- Huber, J., J.W. Payne, and C. Puto 2014. Let's be Honest about the Attraction Effect, *Journal of Marketing Research*, 51/4: 520–5.
- Huber, J., and C. Puto 1983. Market Boundaries and Product Choice: Illustrating Attraction and Substitution Effects, *Journal of Consumer Research*, 10/1: 31–44.
- Jones, K. 1999. Second-Hand Moral Knowledge, *The Journal of Philosophy*, 96/2: 55–78.
- Jones, K., and F. Schroeter 2012. Moral Expertise, *Analyse & Kritik*, 34/2: 217–30.
- Kahneman, D., and G. Klein 2009. Conditions for Intuitive Expertise: A Failure to Disagree, *American Psychologist*, 64/6: 515–26.
- Lewis, M. 2020. A Defense of the Very Idea of Moral Deference Pessimism, *Philosophical Studies*, 177/8: 2323–40.
- Liao, S.M., A. Wiegmann, J. Alexander, and G. Vong 2012. Putting the Trolley in Order: Experimental Philosophy and the Loop Case, *Philosophical Psychology*, 25/5: 661–71.
- Löhr, G. 2019. The Experience Machine and the Expertise Defense, *Philosophical Psychology*, 32/2: 257–73.
- Ludwig, K. 2007. The Epistemology of Thought Experiments: First Person versus Third Person Approaches, *Midwest Studies in Philosophy*, 31/1: 128–59.
- Machery, E. 2012. Expertise and Intuitions About Reference, *Theoria*, 27/1: 37–54.
- Machery, E. 2017. *Philosophy Within Its Proper Bounds*, Oxford: Oxford University Press.
- Machery, E., R. Mallon, S. Nichols, and S. Stich 2004. Semantics, Cross-Cultural Style, *Cognition*, 92/3: 1–12.
- Nado, J. 2014. Philosophical Expertise, *Philosophy Compass*, 9/9: 631–41.
- Nichols, S., and J. Knobe 2007. Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions, *Noûs*, 41/4: 663–85.

- Open Science Collaboration 2015. Estimating the Reproducibility of Psychological Science, *Science*, 349/6251, URL = <<https://doi.org/10.1126/science.aac4716>>
- Palan, S., and C. Schitter 2018. Prolific.ac—A Subject Pool for Online Experiments, *Journal of Behavioral and Experimental Finance*, 17, 22–27.
- Petrinovich, L., and P. O’Neill 1996. Influence of Wording and Framing Effects on Moral Intuitions, *Ethology and Sociobiology*, 17/3: 145–71.
- Ruggeri, K., S. Alí, M.L. Berge, et al. 2019. Not Lost in Translation: Successfully Replicating Prospect Theory in 19 Countries [Preprint], Open Science Framework, URL = <<https://doi.org/10.31219/osf.io/2nyd6>>
- Schindler, S., and P. Saint-Germier. Unpublished Manuscript. *Philosophical Expertise Put to the Test*.
- Schulz, E., E. Cokely, and A. Feltz 2011. Persistent Bias in Expert Judgments About Free Will and Moral Responsibility: A Test of the Expertise Defense, *Consciousness and Cognition*, 20/4: 1722–31.
- Schwitzgebel, E., and F. Cushman 2012. Expertise in Moral Reasoning? Order Effects on Moral Judgment in Professional Philosophers and Non-Philosophers, *Mind & Language*, 27/2: 135–53.
- Schwitzgebel, E., and F. Cushman 2015. Philosophers’ Biased Judgments Persist Despite Training, Expertise and Reflection, *Cognition*, 141, 127–37.
- Sytsma, J., and E. Machery 2010. Two Conceptions of Subjective Experience, *Philosophical Studies*, 151/2: 299–327.
- Thaler, R. 1985. Mental Accounting and Consumer Choice, *Marketing Science*, 4/3: 199–214.
- Thomson, J.J. 1971. A Defense of Abortion, *Philosophy & Public Affairs*, 1/1: 47–66.
- Thomson, J.J. 1985. The Trolley Problem, *The Yale Law Journal*, 94/6: 1395–415.

- Tobia, K., W. Buckwalter, and S. Stich 2013. Moral Intuitions: Are Philosophers Experts?, *Philosophical Psychology*, 26/5: 629–38.
- Tobia, K., G. Chapman, and S. Stich 2013. Cleanliness is Next to Morality, Even for Philosophers, *Journal of Consciousness Studies*, 20/11–12: 195–204.
- Tversky, A., and D. Kahneman 1981. The Framing of Decisions and the Psychology of Choice, *Science*, 211/4481: 453–8.
- Weinberg, J.M. 2007. How to Challenge Intuitions Empirically Without Risking Skepticism, *Midwest Studies in Philosophy*, 31/1: 318–43.
- Weinberg, J.M. 2015. The Methodological Necessity of Experimental Philosophy, *Discipline Filosofiche*, 15/1: 23–42.
- Weinberg, J.M., C. Gonnerman, C. Buckner, and J. Alexander 2010. Are Philosophers Expert Intuiters?, *Philosophical Psychology*, 23/3: 331–55.
- Wiegmann, A., J. Horvath, and K. Meyer 2020. Intuitive Expertise and Irrelevant Options, *Oxford Studies in Experimental Philosophy*, 3: 275–310.
- Wiegmann, A., and M.R. Waldmann 2014. Transfer Effects Between Moral Dilemmas: A Causal Model Theory, *Cognition*, 131/1: 28–43.
- Williamson, T. 2011. Philosophical Expertise and the Burden of Proof, *Metaphilosophy*, 42/3: 215–29.