

Do non-philosophers think epistemic consequentialism is counterintuitive?*

James Andow

March 16, 2016

Abstract

Direct epistemic consequentialism is the idea that X is epistemically permissible iff X maximizes epistemic value. It has received lots of attention in recent years and is widely accepted by philosophers to have counterintuitive implications. There are various reasons one might suspect that the relevant intuitions will not be widely shared among non-philosophers. This paper presents an initial empirical study of ordinary intuitions. The results of two experiments demonstrate that the counterintuitiveness of epistemic consequentialism is more than a philosophers' worry—the folk seem to agree!

1 Introduction

A number of philosophers have recently explored consequentialist accounts in epistemology in some detail. These accounts are modeled on consequentialist theories in normative ethics. A common variety of objection to consequentialist accounts in normative ethics involves highlighting certain supposedly counterintuitive implications that such accounts have. Epistemic consequentialism is no different in this respect. This is a problem for epistemic consequentialists unless they are prepared to bite the bullet or deny the claim that the supposedly counterintuitive implications are indeed counterintuitive.

Philosophers' claims about which things are intuitive can no longer be taken for granted. The advent of experimental philosophy means that the relevant claims can be tested. Are such-and-such supposedly counterintuitive implications actually so very contrary to our ordinary ways of thinking about the relevant philosophical

*This paper is forthcoming in *Synthese*. Please refer to the published version for citations, etc.

quantities or concepts? In debates about consequentialism in ethics, experimental philosophers have made considerable contributions.

This current paper aims to bring similar tools to the debate about epistemic rationality. While the current paper represents only a first step, the results reported suggest that epistemic consequentialism does not accord with ordinary intuitions about epistemic rationality. The paper proceeds as follows. §2 outlines epistemic consequentialism using the more familiar ethical consequentialism as a reference point and outlines the supposedly counterintuitive implications of epistemic consequentialism. §3 briefly describes some experimental work on ethical consequentialism. §4 describes and reports the results of two experiments conducted which examine intuitions about epistemic consequentialism. §5 draws these results together, highlights some shortcomings, and makes some suggestions for the future direction of this research.

2 Consequentialism: Epistemic and Ethical

Consequentialist theories in normative ethics will be familiar to most readers. Perhaps the most familiar is act utilitarianism. According to this account, simply speaking, pleasure is a positive moral value and pain is a negative moral value. It is only morally permissible to do that which, of all the alternative actions available to an agent, maximizes moral utility. The commitments of such a view can be illustrated using an infamous ‘trolley case’ (this version is taken from [Liao et al., 2011](#)).

Switch A runaway trolley is headed toward five innocent people who are on the track and who will be killed unless something is done. Abigail can push a button, which will redirect the trolley onto a second track, saving the five people. However, on this second track is an innocent bystander, who will be killed if the trolley is turned onto this track.

Given the two available options, the act utilitarian thinks the only morally acceptable action is to push the button and switch the tracks.

The question of which actions are justified *morally speaking* has an epistemic analogue. Given my evidence, what am I warranted in believing? Which epistemic doings (e.g., believing, inferring, judging) are epistemically permissible? The epistemic consequentialist holds that ‘the epistemically right (e.g., the justified) is to be understood in terms of conduciveness to the epistemic good (e.g., true belief)’ ([Ahlstrom-Vij and Dunn, 2014](#)). The epistemic equivalent to the act utilitarian, the *direct epistemic consequentialist* would endorse the claim that X is epistemically permissible iff X maximizes epistemic value. Such accounts have received some notable attention in recent years.¹

¹See, e.g., [Berker \(2013\)](#); [Briesen \(forthcoming\)](#); [Carr \(2015\)](#); [Firth \(1981\)](#); [Goldman \(1986\)](#);

As many readers will know, act utilitarianism is often claimed to have counter-intuitive implications. One stark case which has been used to make this point is Push (again this version is taken from [Liao et al., 2011](#)).²

Push A runaway trolley is headed toward five innocent people who are on the track and who will be killed unless something is done. Abigail can push a button, which will activate a moveable platform that will move an innocent bystander in front of the trolley. The runaway trolley would be stopped by hitting the innocent bystander, thereby saving the five but killing the innocent bystander.

The consequences of acting and not acting in both Switch and Push are identical. So the utilitarian should give the same answer in each. Taking the action which leads to the death of only one person rather than five seems to maximize utility in each case. However, critics of utilitarianism claim, the idea that it is morally acceptable to push the button in Push is counterintuitive. At the very least, they claim, there is an intuitive difference between the two cases—one seems less acceptable than the other—a difference which the utilitarian cannot accommodate.³

Direct epistemic consequentialists face similar charges. For example, [Jenkins \(2007\)](#) presents the following case which highlights the relevant counterintuitive implication of a ‘veritist’ version of epistemic consequentialism:⁴

Truth Fairy Suppose you start with no reason to believe that *p* is true and no reason to believe that it is false. The Truth Fairy is a very powerful being, and she makes you the following credible offer: you accept *p* as true, and she will make your epistemic situation very, very good overall. She will arrange for you to have many, many true, justified, knowledgeable beliefs, and very, very few false, unjustified or unknowledgeable ones. However, she does not guarantee that your trust in *p* itself will have any particular epistemic status as a result of her actions.

Jenkins claims that intuitions ‘rebel’ at the thought that trust in *p* would be epistemically permissible as direct epistemic consequentialism would suggest. As [Jenkins and Elstein \(Forthcoming\)](#) put it,

[Greaves \(2013\)](#); [Jenkins and Elstein \(Forthcoming\)](#); [Jenkins \(2007\)](#); [Littlejohn \(2012\)](#); [Maitzen \(1995\)](#); [Percival \(2002\)](#).

²In original presentations such cases, the innocent bystander is stated to be a ‘Fat Man’ (e.g., [Thomson, 1985](#)).

³It should be noted that act utilitarians typically do not contest the claim that such implications are counterintuitive. The strategy rather tends to be to explain away the the relevant intuitions.

⁴Many others have articulated similar counterintuitive implications of epistemic consequentialist accounts. I focus on Jenkins’s articulation here.

Whatever kind of reasons the Truth Fairy might give you to trust in p , she cannot by these means make it epistemically permissible for you to do so.

Note that the Truth Fairy case is, in some respects, the epistemic analogue of Push. The protagonist has an opportunity to maximize (epistemic) utility by doing something which might in other circumstances seem (epistemically) suspect. The question is whether the maximization of utility is the driving concern. There are dis-analogies too, but we will come to those later.

3 Testing Intuitions

When philosophers make points or arguments which rely upon claims about what is and what is not intuitive, they are making empirical claims.⁵ These empirical claims can be tested empirically.⁶

Cases such as Switch and Push have been subject to empirical attention. Indeed the debate has gone far beyond whether or not people agree that acting in the Push case is morally acceptable.⁷ Nonetheless, the results demonstrate that the act utilitarian has some work to do; for, acting in cases such as Push is judged to be rather less morally acceptable than in cases such as Switch. There are various types of response open to the utilitarian. For instance, they might choose to bite the bullet or attempt to explain away these intuitions (to give some reason for thinking that it is okay that their account has these counterintuitive implications about such cases). However, the empirical result demonstrate that the utilitarian has a heavy counterintuitive burden which must be dealt with in some way.

It is worth noting another way that this debate in ethics might have gone. Had the empirical results been different—had people been just as willing to sacrifice the innocent bystander in Switch and Push—the result would have been a great relief to the consequentialist. The results would have cleared utilitarians of a supposed large counterintuitive burden. There'd be no bullet biting involved in sticking to their guns, no pesky misfiring intuitions to explain away.

⁵I shall assume that the relevant points do rely upon such claims. However, some have recently questioned the role of intuitive judgments in the use of hypothetical cases such as these in philosophy (notably Cappelen, 2012; Deutsch, 2015).

⁶Testing philosophers' empirical claims is one aspect of the aims of experimental philosophy. For various perspectives on this developing field, see, Alexander (2012); Andow (Forthcoming); Feltz (2009); Guttenplan (2011); Kauppinen (2007); Knobe (2007a,b, 2010); Knobe et al. (2011); Knobe and Nichols (2008); Levin (2009); Machery and O'Neill (2014).

⁷See, e.g., Ahlenius and Tännsjö (2012); Fischer and Ravizza (1992); Greene (2008); Kahane (2013); Liao et al. (2011); Lombrozo (2009); Petrinovich and O'Neill (1996); Petrinovich et al. (1993).

The question this paper addresses could then be framed as follows: do epistemic consequentialists truly face the counterintuitive burden which they are alleged to face? One might wonder why this is a question worth asking. Is there any reason to think there might be a mismatch between the intuitions of ordinary folks and philosophers in this particular case? I think there are at least two reasons:

1. Various elements in ethical trolley cases are absent from the relevant epistemic cases (in ways that I take to diminish the likelihood of a ‘deontological’ response in the epistemic case). For example:
 - (a) The high affect nature of Push. The idea of ordinary folks having a real gut reaction, in the sense of one really *felt* in the pit of the stomach, about epistemic cases like the Truth Fairy case seems somewhat fanciful. So, it seems more likely that folks would treat epistemic equivalents of Push and Switch in similar ways.
 - (b) The high cost of not acting in Push. However epistemic value and moral value stack up, I take it no one is seriously going to entertain the idea that a false belief or potentially false belief is as bad as a death by runaway train. Consequently, folks may have few additional qualms in the epistemic equivalent of Push, despite drawing an important line when it comes to using someone in a way that kills them in order to save five lives.
 - (c) An analogue of the idea of an ‘innocent bystander’. In the moral case, this notion of ‘innocence’ seems to bestow the bystander with a certain moral purity/sanctity/to-be-protectedness. In the relevant epistemic cases, there doesn’t seem to be anything equivalent to this ‘innocence’, e.g., some to-be-protected True Belief which would be destroyed in the event of the protagonist pushing the button.
2. There is reason to doubt that people will entertain any strong prohibitions when it comes to the epistemic domain. I take this consideration to count in favor of the hypothesis that insofar as folks have any intuitions they will not be one’s which put pressure on consequentialism. Why? In absence of strong intuitive prohibitions, I think that:
 - (a) The default response to any case would be that the action is permissible. One might object at this point that the maximizing consequentialist doesn’t think it would be okay, e.g., to not believe in the Truth Fairy case. However, if ordinary folks had no strong intuitions about epistemic permissibility, then (i) the consequentialist would have much less of a problem than if ordinary folks had strong anti-consequentialist intuitions as its opponents allege; failing to have intuitive support is a rather

better position than having strong intuitive opposition, and (ii) moreover, the problem the consequentialist would face would be one shared equally with the opponent who thinks one should not take the Truth Fairy bargain.

- (b) If forced to make a choice, e.g., between believing and not believing in the Truth Fairy case, it would be reasonable to expect folks to utilize some kind of simple cost-benefit analysis. The only salient costs and benefits in something like a Truth Fairy case are the truth and falsity of the relevant claims/beliefs and so a simple cost-benefit analysis, absent any strong intuitive prohibitions, might be expected to favor the consequentialist option.

In other words, it is tempting to think that the unacceptability of believing, in the Truth Fairy case, is something of a ‘philosopher’s reaction’ and thus the relevant resistance to epistemic consequentialism a something of a ‘philosopher’s worry’.

So, there is specific reason to think it is worth testing philosophers’ claims about intuitions about epistemic rationality. However, the issues I have just raised also draw attention to some difficulties which the current project faces. One advantage in the project of testing *moral* intuitions about the permissibility of actions in trolley cases is that the cases (e.g., Push or its predecessors) and the questions are comparatively comprehensible by the average participant in an experiment. In the current project, testing epistemic intuitions about the epistemic permissibility of such things (e.g., taking the Truth Fairy’s bargain), we do not have that advantage. The Truth Fairy case is rather less down to earth and it employs concepts which we cannot expect the average participant to understand (e.g., ‘epistemic status’). One solution would be to recruit only participants who were familiar with such concepts.⁸ However, given that such a sample would primarily consist of participants with at least philosophy undergraduate degrees, such an experiment would have limited claim on *testing ordinary pretheoretic judgments about such cases*.⁹

4 Two Experiments

In this section, I report two experiments which represent a first step in testing intuitions about epistemic consequentialism. The design attempts to overcome some of the difficulties discussed in the previous section.

⁸Let me clarify. What I think the average participant is unfamiliar with is primarily the philosophical vocabulary, e.g., the word ‘epistemic’. I do think they are familiar with the concept in the sense that they distinguish between epistemically permissible and impermissible beliefs.

⁹I shall assume that it is such judgments which are relevant to the project. However, I acknowledge that this is likely not completely uncontroversial. Nonetheless, I won’t be defending my stance.

4.1 Experiment 1

In a first experiment, I attempted to see whether participants saw an intuitive difference between cases in which a protagonist can maximize utility by some neutral means or (like in the Truth Fairy case) maximize utility by some means which might otherwise be judged epistemically dubious (e.g., accepting a claim without any evidence).

Participants 80 participants were recruited using Prolific Academic.¹⁰ All participants were resident in the UK, native English speakers, and at least 18 years old. Mean age was 30 years old. 35 participants identified as Male. Participants were rewarded for their participation (£0.25 @ £7.50/h based on a predicted completion time of 2 minutes).

Methods The design was between-subjects with each participant being assigned to one of four conditions. In each condition, participants were presented with a hypothetical scenario involving a researcher named Sarah. The main contrast of interest is that between the following two scenarios.

- (a) **Equipment** Sarah is conducting some research. She is investigating 5 key claims. Given the current direction of her research, her results will be wrong about these 5 key claims. Sarah could make a particular change in the direction of her research. She could use some different equipment. By making this change, she could ensure that her results are correct about 4 of the 5 key claims. However, her results would still be wrong about the remaining claim.
- (b) **False Assumption** Sarah is conducting some research. She is investigating 5 key claims. Given the current direction of her research, her results will be wrong about these 5 key claims. Sarah could make a particular change in the direction of her research. She could make a particular assumption about 1 of the 5 claims and use that assumption in her research. By making this change, she could ensure that her results are correct about the other 4 claims. However, this assumption about the remaining claim happens to be false.

This second scenario is similar to the Truth Fairy cases discussed above. Sarah can maximize utility by making an assumption without evidence which happens to be false (something which in other circumstances would generally not be okay). By stating that the epistemic state used to maximize utility ‘happens to be false’, however, False Assumption departs from the structure of the Truth Fairy case. In a

¹⁰Prolific Academic is a UK-based alternative to Amazon M-Turk. See <https://prolificacademic.co.uk/>.

Truth Fairy case, the accepted claim is not stated to be false, rather it is stated, e.g., that the protagonist has no fairy-independent epistemic reason either for accepting or rejecting the claim in question. In order to explore whether truth/falsity of the relevant claim made a difference to participant's intuitions, two further cases were used.

(c) **True Assumption** Sarah is conducting some research She is investigating 5 key claims. Given the current direction of her research, her results will be wrong about these 5 key claims. Sarah could make a particular change in the direction of her research. She could make a particular assumption about 1 of the 5 claims and use that assumption in her research. By making this change, she could ensure that her results are correct about the other 4 claims. Moreover, this assumption about the remaining claim happens to be true.

(d) **No Reason Assumption** Sarah is conducting some research. She is investigating 5 key claims. Given the current direction of her research, her results will be wrong about these 5 key claims. Sarah could make a particular change in the direction of her research. She could make a particular assumption about 1 of the 5 claims and use that assumption in her research. By making this change, she could ensure that her results are correct about the other 4 claims. However, she has no evidence for or against the assumption, she has no particular reason either for thinking it is true or for thinking it is false.

In each condition, the main case was followed by this claim in order to guard against the potential complications introduced if participants thought Sarah was simply a poor researcher:

Whether or not Sarah changes her research, any false results she obtains are not due to any mistake in her research. It is just one of those things, research doesn't always produce accurate results.

Then participants answered a single question on a sliding scale from 0 (completely disagree) to 100 (completely agree). The starting position of the slider was 50. The wording of the statement was

'It is okay for Sarah to change her research by (a) using some different equipment / (b) making an assumption which happens to be false / (c) making an assumption which happens to be true / (d) making an assumption for which she has no evidence.'

Finally, some basic demographic information was collected.

Table 1: Median and means by scenario

	Median	Mean
(a) Equipment	71	68.47
(b) False Assumption	18.5	26.9
(c) True Assumption	32	41
(d) No Reason Assumption	35.5	38.9

Table 2: p-values and effect sizes for pairwise comparisons of medians

	(b)	(c)	(d)
(a)	$p = .001^{**}, \phi = .546$	$p = .003^{**}, \phi = .47$	$p = .023^*, \phi = .363$
(b)	-	$p = .585, \phi = .085$	$p = .197, \phi = .204$
(c)	-	-	$p = .443, \phi = .120$

Results A clear pattern is suggested by the median and mean scores for each scenario (shown in Table 1): agreement in the Equipment case and disagreement in the other cases. The scores for the four scenarios were compared using Mood’s median test.¹¹ There was a significant omnibus effect, $\chi^2(3, N = 80) = 13.285, p = .004, \phi = .408$. Post-hoc chi-square tests revealed that a greater proportion of participants in the Equipment condition (a) gave a score at or above the median than in each of the other scenarios (b, c, d) (see Table 2 for effect sizes and p-values). There was no significant difference between the three ‘assumption’ conditions (b, c, d).

Discussion It seems clear that, while participants are in general happy with making changes to maximize epistemic utility, they are not happy with those changes involving the use of an assumption for which the protagonist has no evidence. This suggests that our participants have intuitions clearly in line with the Truth Fairy type cases and that maximizing epistemic utility are not the only grounds upon which they judge the acceptability of the doings of epistemic agents. Note also that whether the assumption happens to be true, happens to be false, or is not stated, seems to have no effect on participants’ intuitions. Were there a large (veritist) consequentialist component to their intuitions about these cases, one would expect the truth/falsity of the assumption to have an effect.

¹¹This is a close non-parametric equivalent to a one-way ANOVA and can be used to compare more than two groups (unlike, e.g., Man-Whitney’s U test). Non-parametric tests were used as (a) visual inspection suggests a bimodal distribution for some scenarios, and (b) Shapiro-Wilk tests confirm that the data can’t be assumed to be normally distributed (scenario a, $W(19)=.881, p=.023$; scenario, b $W(20)=.842, p=.004$; scenario c, $W(21)=.938, p=.201$; scenario d, $W(20)=.897, p=.036$).

4.2 Experiment 2

In a second experiment, I looked at another class of cases. These are designed to be more directly analogous to Switch and Push (the moral cases discussed earlier). The cases in experiment 1 are not directly analogous for an important reason: they are not really cases in which the protagonist has an opportunity to maximize utility by bringing about some negative outcome which would not have come about anyway. For example, in Equipment, Sarah was going to be wrong about the claim which she ends up being wrong about regardless of whether she makes the change.

Because of this feature of Equipment, the decision Sarah faces is not whether to choose one negative thing in order to prevent a *different* more negative thing. This is important since a critic might argue that the difference observed in experiment 1 would go away if presented with versions of the equipment scenario and assumption scenarios in which Sarah did have to make such a decision. So, the second experiment, puts this thought to the test.

Participants 60 participants were recruited using Prolific Academic. Participants who had already taken Experiment 1 were prevented from taking part. All participants were resident in the UK, native English speakers, and at least 18 years old. Mean age was 29.7 years old. 26 participants identified as Male. Participants were rewarded for their participation (£0.25 @ £7.50/h based on a predicted completion time of 2 minutes).

Methods The design was between subjects with each participant being assigned to one of three conditions. Again, in each condition, participants were presented with a hypothetical scenario involving Sarah a researcher. This time the main contrast is between two cases. This time the cases are stricter analogues to the simple switch and push trolley cases.

(e) **Equipment-switch** Sarah is conducting some research. She is investigating 5 key claims: A, B, C, D, and E. Given the current direction of her research, her results will be correct about 1 of the 5 key claims (A), but incorrect about the remaining 4 (B, C, D, and E).

Sarah could make a particular change in the direction of her research. She could use some different equipment. By making this change, she could ensure that her results are correct about 4 of the 5 key claims (B, C, D, and E). However, her results about the other claim (A) would be incorrect.

(f) **False Assumption-push** Sarah is conducting some research. She is investigating 5 key claims: A, B, C, D, and E. Given the current direction of her research,

her results will be correct about 1 of the 5 key claims (A), but incorrect about the remaining 4 (B, C, D, and E).

Sarah could make a particular change in the direction of her research. She could make a particular assumption about 1 of the 5 claims (A) and use that assumption in her research. By making this change, she could ensure that her results are correct about 4 of the 5 key claims (B, C, D, and E). However, her assumption about the other claim (A) happens to be incorrect.

A third scenario was also used to continue to allow comparisons with intuitions about more Truth Fairy type cases in which the valence of the used epistemic state is unknown.

(g) No Reason Assumption-fairy Sarah is conducting some research. She is investigating 5 key claims: A, B, C, D, and E. Given the current direction of her research, her results will be correct about 1 of the 5 key claims (A), but incorrect about the remaining 4 (B, C, D, and E).

Sarah could make a particular change in the direction of her research. She could make a particular assumption about 1 of the 5 claims (A) and use that assumption in her research. By making this change, she could ensure that her results are correct about 4 of the 5 key claims (B, C, D, and E). However, she has no evidence for or against the other claim (A), she has no particular reason either for thinking it is true or for thinking it is false.

In each condition, the passage was followed by the same qualification as in Experiment 1, the probes were the same, and demographic data were collected.

Results The median and mean scores for each scenario are shown in [Table 3](#). The same clear pattern is suggested.¹² Mood's median test revealed a significant omnibus effect, $\chi^2(2, N = 60) = 21.758, p < .0005, \phi = .602$. Post-hoc chi-square tests revealed medium to large significant differences between Equipment and each other condition (see [Table 4](#) for details). There were no significant differences between the two assumption conditions.

Discussion It seems that the difference observed in Experiment 1 is not best attributed to the fact that the equipment scenario did not involve Sarah bringing about a new negative consequence which would not otherwise have happened. The scenarios used in this second experiment do not have this feature—and are so more analogous to the Switch and Push trolley cases—but we still find similar results.

¹²Again, Shapiro-Wilk tests confirm that the data can't be assumed to be normally distributed (scenario e, $W(20) = .843, p = .004$; scenario f, $W(20) = .860, p = .008$; scenario g, $W(20) = .832, p = .003$). So, again, non-parametric tests are used.

Table 3: Median and means by scenario

Scenario	Mean	Median
(e) Equipment-switch	72.4	84.5
(f) False Assumption-push	32.3	21
(g) No Reason Assumption-fairy	25.35	20

Table 4: p-values and effect sizes for pairwise comparisons of medians

	(f)	(g)
(e)	$p < .0005^{**}$, $\phi = .568$,	$p < .0005^{**}$, $\phi = .704$
(f)	–	$p = .288$, $\phi = .168$

Again, note that whether the assumption is stated to be false seems to have no effect on participants’ intuitions (a result which wouldn’t fit well with the idea that participants are veritist epistemic consequentialists).

5 Does epistemic consequentialism have counterintuitive implications?

While further work is needed, I tentatively conclude that epistemic consequentialism seems to have consequences which are counterintuitive to the folk and not simply to philosophers. Direct epistemic consequentialism says that X is epistemically permissible iff it maximizes epistemic utility. In the assumption cases, Sarah’s assumption should be judged epistemically permissible by such standards. My results suggest that this consequence is clearly counterintuitive. If correct, these results mean that the supposed counterintuitiveness of epistemic consequentialism is not just a ‘philosophers’ worry’ as one might have thought.¹³

I should emphasize that my conclusions are made only *tentatively* in light of the small sample size and the nature of the probe used in both Experiment 1 and 2. In order to avoid asking participants to judge whether the relevant epistemic doings were ‘warranted’ or ‘epistemically permissible’—terms which one shouldn’t expect participants to understand—the main probe used asks simply whether the relevant

¹³Of course, the epistemic consequentialist may find some way to explain away such intuitions, be otherwise prepared to bite the bullet, or adapt their account in order to avoid these implications (see, e.g., the attempt in [Jenkins and Elstein, Forthcoming](#)). One potential way to accommodate the results might be to suggest that participants are operating with a veritist version of epistemic consequentialism which assigns a greater negative value to false beliefs (or judgements/assumptions) than positive value to true beliefs, e.g., false beliefs are more than five times worse than true beliefs are good. It would be interesting to explore this possibility further.

epistemic doings are ‘okay’. I think participants’ responses likely do reflect their thoughts on the epistemic permissibility of the relevant epistemic doings. However, this isn’t the only available interpretation. It is possible that the probe used may simply tap participants’ moral attitudes. For example, the results *could* be explained by participants judging the use of different equipment to be morally permissible but the making of assumptions without independent reason to be morally impermissible.¹⁴

However, I don’t want to be too concessive. I’ll be the first to admit that the phrasing adopted in this paper is likely not the final word in how best to ask participants about epistemic permissibility. This is a tricky issue and further work will hopefully find even better solutions. Nonetheless, I think the current results can be taken to be suggestive since, on the face of it, it seems rather unlikely that participants are reporting moral attitudes. I say this because it is unlikely that participants find the act of making an assumption to be *morally* objectionable. Even in the cases in which Sarah’s possible assumption turns out to be false, the scenarios give no sense that Sarah is engaged in any clearly morally objectionable act of deception or misleading.¹⁵ So the take home message should be that (tentatively) this initial study of intuitions about epistemic rationality shows that epistemic consequentialism really does run contrary to ordinary folks intuitions—not just to those of philosophers.

Acknowledgements

Thanks to Aimie Hope, Pendaran Roberts, and an audience in Nottingham whose comments at various stages of this project were very helpful.

References

- Ahlenius, H. and Tännsjö, T. (2012). Chinese and westerners respond differently to the trolley dilemmas, *Journal of Cognition and Culture* 12(3-4): 195–201.
- Ahlstrom-Vij, K. and Dunn, J. (2014). A defence of epistemic consequentialism, *The Philosophical Quarterly* 64(257): 541–551.
URL: <http://pq.oxfordjournals.org/content/64/257/541.abstract>

¹⁴Another possible concern is that participants are epistemic consequentialists but are not veritist, e.g., they think that the relevant epistemic value is knowledge rather than true believe, or endorse a different version of consequentialism, e.g., an indirect form or one which involves the maximisation of *expected* epistemic value. The results reported here may be compatible with a number of those hypotheses. Further work will be needed to test them.

¹⁵Note that if it were the case that the current results reflected a difference in moral intuitions, the results would be all the more striking. It really would be striking if participants moralize assumption-making in this way. For, one would typically think that assumptions are relatively morally neutral in research unless they have morally-objectionable content, e.g., racist or sexist assumptions.

- Alexander, J. (2012). *Experimental Philosophy: An Introduction*, Polity, Cambridge.
- Andow, J. (Forthcoming). Critical notice of booth and rowbottom (eds) ‘intuitions’, *Analysis* .
- Berker, S. (2013). The rejection of epistemic consequentialism, *Philosophical Issues* 23(1): 363–387.
- Briesen, J. (forthcoming). Epistemic consequentialism: Its relation to ethical consequentialism and the truth-indication principle, in P. Schmechtig and M. Grajner (eds), *Epistemic Reasons, Norms, and Goals*.
- Cappelen, H. (2012). *Philosophy Without Intuitions*, OUP.
- Carr, J. (2015). Epistemic utility and the aim of belief, *Work in progress* .
- Deutsch, M. (2015). *The Myth of the Intuitive*, MIT Press.
- Feltz, A. (2009). Experimental philosophy, *Analyse and Kritik* 31: 201–219.
- Firth, R. (1981). Epistemic merit, intrinsic and instrumental, *Proceedings and Addresses of the American Philosophical Association* 55(1): 5–23.
- Fischer, J. M. and Ravizza, M. (1992). *Ethics : problems and principles*, Harcourt Brace Jovanovich College Publishers, Fort Worth.
- Goldman, A. (1986). *Epistemology and Cognition*, Harvard University Press, Cambridge, MA.
- Greaves, H. (2013). Epistemic decision theory, *Mind* 122(488): 915–952.
- Greene, J. D. (2008). *The secret joke of Kant’s soul*, Cambridge, MA, US: MIT Press, pp. 35–80.
- Guttenplan, S. (2011). Experimental philosophy, *Mind & Language* 26(4): 452–452.
- Jenkins, C. I. and Elstein, D. (Forthcoming). The truth fairy and epistemic consequentialism, in N. J. L. L. Pedersen and P. Graham (eds), *Epistemic Entitlement*, Oxford University Press.
- Jenkins, C. S. (2007). Entitlement and rationality, *Synthese* 157(1): 25–45.
- Kahane, G. (2013). The armchair and the trolley: an argument for experimental ethics, *Philosophical Studies* 162(2): 421–445.
- Kauppinen, A. (2007). The rise and fall of experimental philosophy, *Philosophical Explorations* 10(2): 95–118.

- Klein, G. (2015). A naturalistic decision making perspective on studying intuitive decision making, *Journal of Applied Research in Memory and Cognition* .
- Knobe, J. (2007a). Experimental philosophy, *Philosophy Compass* 2(1): 81–92.
- Knobe, J. (2007b). Experimental philosophy and philosophical significance, *Philosophical Explorations* 10: 119–122.
- Knobe, J. (2010). Experimental philosophy and the analytic project.
- Knobe, J., Buckwalter, W., Robbins, P., Sarkissian, H., Sommers, T. and Nichols, S. (2011). Experimental philosophy, *Annual Review of Psychology* 63.
- Knobe, J. and Nichols, S. (2008). *An Experimental Philosophy Manifesto*, Oxford University Press, Oxford, chapter 1, pp. 3–14.
- Levin, J. (2009). Experimental philosophy, *Analysis* 69(4): 761–769.
- Liao, S. M., Wiegmann, A., Alexander, J. and Vong, G. (2011). Putting the trolley in order: Experimental philosophy and the loop case, *Philosophical Psychology* 25(5): 661–671.
- Littlejohn, C. (2012). *Justification and the Truth Connection*, Cambridge University Press.
- Lombrozo, T. (2009). The role of moral commitments in moral judgment, *Cognitive Science* 33(2): 273–286.
- Machery, E. and O’Neill, E. (2014). *Current Controversies in Experimental Philosophy*, Routledge.
- Maitzen, S. (1995). Our errant epistemic aim, *Philosophy and Phenomenological Research* 55(4): 869–876.
- Percival, P. (2002). Epistemic consequentialism: Philip percival, *Aristotelian Society Supplementary Volume* 76(1): 121–151.
URL: <http://dx.doi.org/10.1111/1467-8349.00092>
- Petrinovich, L. and O’Neill, P. (1996). Influence of wording and framing effects on moral intuitions, *Ethology and Sociobiology* 17(3): 145–171.
- Petrinovich, L., O’Neill, P. and Jorgensen, M. (1993). An empirical study of moral intuitions: Toward an evolutionary ethics, *Journal of Personality and Social Psychology* 64(3): 467–478.
- Thomson, J. J. (1985). The trolley problem, *The Yale Law Journal* 94(6): 1395–1415.