

Lay intuitions about epistemic normativity

Pendaran Roberts¹ · James Andow² ·
Kelly Ann Schmidtke¹

Received: 20 July 2016 / Accepted: 8 March 2017
© The Author(s) 2017. This article is an open access publication

Abstract Recent empirical work on non-philosophers’ intuitions about epistemic normativity reveals patterns that cannot be fully accounted for by direct epistemic consequentialism. On the basis of these results, one might picture participants as “epistemic deontologists.” We present the results of two new experiments that support a more nuanced picture. We examine intuitions about guesses and hypotheses, and about beliefs. Our results suggest a two-factor model of intuitions, wherein both consequentialist and non-consequentialist considerations affect participants’ judgments about epistemic permissibility.

1 Epistemic normativity and consequentialism

Epistemically speaking, is it okay for me to believe P, or not? This is a question about epistemic normativity. In recent years, the idea that epistemic normativity should be understood in consequentialist terms has received quite a bit of attention.¹ The basic idea behind epistemic consequentialism is that the epistemic right or permissible is to be understood fully in terms of conduciveness to the epistemic good. [Berker \(2013\)](#)

¹ See, for example, the forthcoming volume on the subject (Ahlstrom-Vij and Dunn forthcoming), as well as [Ahlstrom-Vij and Dunn \(2014\)](#), [Berker \(2013\)](#), [Goldman \(2015\)](#), [Hedberg \(2014\)](#), [Jenkins \(2007\)](#), [Klausen \(2009\)](#), and [Percival \(2002\)](#).

✉ Pendaran Roberts
Pendaran@mac.com

¹ University of Warwick, Coventry, UK

² University of Reading, Reading, UK

provides a framework for thinking about consequentialist theories.² According to this framework, theories can vary along three dimensions: final value, overall value, and deontic theory.

A theory of final value says what is ultimately or fundamentally of epistemic value or disvalue. For example, some of the most often discussed versions of epistemic consequentialism say that the only things of final value are true beliefs and the only things of final disvalue are false beliefs. We will follow the convention of using “veritist” to refer to those theories that endorse such an account of final value. It is important to note that veritism is not the only theory of final value that is available to the epistemic consequentialist. There are also other understandings of epistemic final value available that take into account other features of beliefs (e.g., how important, informative, or interesting they are) or belief sets (e.g., coherence) in addition to or in place of purely veritist concerns.³

A theory of overall value provides a way of ranking epistemic acts like the forming of beliefs. For example, a veritist might say that the overall value of forming any particular belief is the net balance of true beliefs over false beliefs that would be brought about by forming that belief. An alternative veritist account might say that the overall value of a belief is the net balance of true over false beliefs that the agent *expects* would be brought about by forming that belief. An account of overall epistemic value thus provides a way to rank the options available to an agent. It might tell one, for example, that belief *x* is better than belief *y* which is better than belief *z*. Epistemic consequentialists who accept a richer account of final value, in which something other than truth or falsity is of final value, must adapt their theory of overall value accordingly. Note that one might want to advance a theory of overall value that assigns overall values (and thus a ranking) to something other than beliefs, for instance, to other epistemic “actions” such as assumptions, hypotheses, or belief-forming processes.

A deontic theory delineates the epistemic right and wrong.⁴ A theory of overall value does not itself tell us which epistemic acts are permissible or impermissible: the fact that *x* is better than *y* which is better than *z* is compatible with all three acts being permissible or impermissible. One straightforward deontic theory, especially relevant to this paper, is direct epistemic consequentialism (DEC): it is right for a subject to *x* at a specific time if and only if *x* maximizes overall epistemic value for that subject at that time.⁵ However, as is the case in normative ethics, an epistemic consequentialist has a great variety of deontic theories to choose from; these choices include indirect

² A similar framework is taken up by [Dunn \(2016\)](#).

³ One might also construct a consequentialist view where it is not beliefs which are of ultimate value/disvalue but, for example, subjective probability assignments or knowledge. See [Dunn \(2016\)](#) for a survey of various views that might fall under the categorization of “epistemic consequentialist.”

⁴ “Deontic” here means something like “concerning the right” and should not be confused with “deontology.”

⁵ DEC represents a simple application of a maximizing consequentialist norm to the permissibility of epistemic acts. The distinction between “direct” and “indirect” forms of epistemic consequentialism and the analogy with normative ethics are explored in some detail in [Percival \(2002\)](#).

consequentialist theories, satisficing accounts, and so on.⁶ We will encounter some of these in a little more detail later on.

2 The role of intuitions

Epistemic consequentialism has some attractive characteristics (Dunn 2016). However, questions have been raised about certain counterintuitive implications of epistemic consequentialist views like DEC (Berker 2013; Elstein and Jenkins forthcoming). Although such considerations may not have the final word, if a view has counterintuitive implications, this is often seen by epistemologists as a strike against it. One of the best ways to bring out the potentially damaging counterintuitive implications is by looking at some thought experiments: the Truth Fairy case (based on Elstein and Jenkins forthcoming) and the John Doe case (Berker 2013, p. 369; Firth 1981).

Truth Fairy. Suppose you start with no reason to believe that p is true and no reason to believe that it is false. Now suppose you come across the Truth Fairy who makes you the following offer: if you believe p , she will make your epistemic situation very, very good overall. She will arrange for you to have whatever your preferred account of epistemic value says is very, very good—many, many true beliefs, and very, very few false ones, a completely coherent and maximally informative belief set, etc. However, she does not guarantee that your trust in p itself will have any particular epistemic status as a result of her actions.

John Doe. John Doe is a brilliant set theorist who is on the cusp of proving the Continuum Hypothesis: all he needs is six more months. But, alas, poor John is suffering from a serious illness that, according to his doctors, will almost certainly kill him in two months' time. John stubbornly clings to a belief that he will recover from his illness, and not only does this belief comfort him, but—let us suppose—it in fact significantly raises the chances that he will live for the six months that he needs both to complete his proof and to derive from it a variety of consequences for the rest of set theory. In other words, John's belief that he will recover is a causal means to his procuring a large number of true set-theoretic beliefs sometime in the future.

These thought experiments provide the foundation of an intuition based argument against certain forms of epistemic consequentialism (Berker 2013; Elstein and Jenkins forthcoming). DEC, for instance, is committed to the idea that it is epistemically permissible, indeed required, to accept the Truth Fairy's offer and for John Doe to believe that he will recover from his illness (supposing a veritist view of final value). However, based on thought experiments such as those mentioned, one can see that intuitively it is not *epistemically* acceptable to believe something that one has no independent reason to believe simply because that is what will bring about the best

⁶ See Percival (2002) for some details on the variety of accounts available. Notably, reliabilism in epistemology can be characterized as a non-maximizing veritist form of indirect consequentialism about doxastic justification. This is articulated nicely by Berker (2013) who draws on Firth (1978, 1981), although see also Dunn and Ahlstrom-Vij (forthcoming).

epistemic consequences. This means, concludes the argument, that certain forms of consequentialism are problematic. In other words, if consequentialism were right, consequences would be all that mattered, but something else intuitively matters in cases like these, as demonstrated by the thought experiments.⁷

3 Lay intuitions

This article is concerned with folk epistemology and with examining folk intuitions about justification (regarding guesses, hypotheses, and beliefs). If intuitions are relevant to philosophical theories, as the Truth Fairy and John Doe thought experiments presuppose, as well as the intuition based argument, there is reason to be concerned with what non-philosophers' intuitions are (hereafter "lay intuitions"). This is not to say that lay intuitions are preferable to the intuitions of philosophers, but neither is it to deny this. It is only to say that if we are concerned with what is intuitive, it is worth investigating what non-philosophers think, and for a few reasons:

1. Lay intuitions may be less prone to theoretical bias and so better indicators of what is genuinely intuitive than philosophers' intuitions. (Goldman 2007)
2. Lay intuitions may provide better information about concepts that play a role in our everyday discourse and practice. (Alexander and Weinberg 2007)
3. If lay intuitions and those of philosophers differed, this might raise important methodological questions about the practice of relying on intuitions in this debate. (see, e.g., the discussion in Pritchard 2014)

With regard to 1–3, it is of course important to note the existence of entrenched debates in philosophy about the relative importance of philosophers' intuitions versus lay intuitions (Nado 2014). This is not the place to settle those debates. To whatever extent lay intuitions are relevant to philosophical theories, this paper is relevant to philosophical theories.

Lay intuitions are not an a priori matter but an empirical matter to be investigated using methods similar to those used in the social sciences. Recently, some progress has been made in investigating lay intuitions about epistemic consequentialism (Andow 2016a). Andow's findings lend some support to the claim that lay intuitions are consistent with certain philosophers' intuitions to the effect that it is not epistemically acceptable to believe something that one has no independent reason to believe simply because doing so will increase the positive epistemic consequences. Across two experiments, Andow shows that participants' intuitive judgements about the permissibility of adopting attitudes are sensitive to factors that views like DEC would not predict.

Specifically, participants' permissibility judgements in Andow's experiments were sensitive to the means by which an epistemic good was obtained. In his experiments, some participants were asked to consider cases in which a researcher has the opportunity to change the course of her research by using different equipment that will have the effect of increasing the number of correct results she gets. The participants were then

⁷ It may be possible for the consequentialist to avoid this version of the worry by adopting an indirect form of consequentialism (Elstein and Jenkins forthcoming), although similar worries may then surface elsewhere (see discussions in Berker 2013 and Andow forthcoming).

asked to say whether this change is “okay” on a 100-point scale (from completely disagree to completely agree). The mean responses were 69–72. Other participants were asked to consider cases in which a researcher can increase the number of correct results by making an assumption for which it is never stated that there is any independent motivation and sometimes even stated that there is not. Participants were then asked to say whether this change is okay. The mean responses were 25–41. In summary, even though both changing the equipment and making the assumption increased the epistemic good (the number of correct results), participants found a change of equipment more permissible than changing the research by making an assumption.

Andow’s (2016a) experiments leave some important questions unanswered. Andow sets up his experiments with reference to ethical trolley cases.⁸ This setup invites a particular interpretation about what considerations participants are sensitive to. On the basis of his results, we might picture participants as “epistemic deontologists” who find it permissible to bring about the maximization of epistemic value so long as it does not involve using epistemic disvalue (making an assumption on the basis of no evidence) as a *means* to maximize epistemic utility. However, Andow’s results as they stand do not permit us to make this inference. Both of his experiments involve contrasts between cases which do not (or barely) differ in terms of consequences. So, we cannot conclude that his participants are deontologists who think that consequences make no difference. In other words, to test whether people are deontologists it is necessary to use a factorial design where both consequences and means are varied, but Andow only manipulated the means. Are people epistemic deontologists? We designed and ran two studies to find out.

4 Experiment 1

Experiment 1 was performed to investigate people’s intuitions about the epistemic permissibility of certain epistemic acts (making a guess or forming a hypothesis) performed by a hypothetical individual named Sam. Sam’s decisions were either based on evidence or not, and either the consequences of the decisions were good or bad. For each case, participants were asked whether the act is justified. If participants are consequentialists, we should expect only consequences to affect whether they agree the act is justified. If they are deontologists, we should expect only evidence to affect whether they agree the act is justified. We find that participants do not fit into either category.

4.1 Participants

207 participants recruited via Amazon Mechanical Turk completed the survey. We required participants to have at least a 95% HIT approval rate and at least a US high school diploma. The experiment was run until at least 40 participants in each group answered all the comprehension items correctly (the comprehension items will be explained in “Materials and procedure”). Of the 207 participants, 32 were removed

⁸ The use of trolley cases in moral philosophy originates with Foot (1967) and Thomson (1971, 1976). Trolley cases have subsequently played a big role in experimental philosophy (e.g., Greene 2016).

for answering one or more comprehension questions incorrectly, and 3 were removed for having any formal academic training in philosophy whatsoever. Of the remaining 172 participants, 106 identified as female and the median age was 37.

4.2 Materials and procedure

The survey was created and administered using Qualtrics, Version 2016 and Amazon Mechanical Turk.

Participants only received one Case Type selected at random (Evidence No-Bad, Evidence No-Good, Evidence Yes-Bad, Evidence Yes-Good). The Case Types were designed so that Sam's potential decision would either be based on a guess or on an empirically informed hypothesis, and the consequence would be either bad or good. Participants experienced four Case Versions of their assigned Case Type, presented in a random order. The four Case Versions were designed to be structurally similar. The Case Versions can be found in Table 1.

Each case was presented over two screens. The initial screen presented the selected case with three comprehension items, where two items were true and either the first or last was false. For example, for the Experiment Case Version the first comprehension item read, "Given the current direction of Sam's research, Sam will end up being [wrong or correct] about the 5 key claims." The second item read, "To change the direction of the research, Sam could make a particular [guess or hypothesis]." The third item read, "By changing the direction of the research, Sam will end up being [wrong or correct] about the 5 key claims." After responding to the comprehension items, participants clicked a button to continue.

The second screen presented participants with the same case along with a target item called "Justified" that read, "The [guess or hypothesis] is justified." Participants indicated how much they agree with this item on a sliding scale from 0 to 100 (from completely disagree to completely agree) with the starting position set at 50.

After completing the four Case Versions, participants were asked about their gender, age, experience in formal philosophical academic training, and for any other comments they had about the survey.

4.3 Results

Means and standard deviations for all cases are displayed in Table 2. Descriptively, the Evidence No-Bad cases garnered the lowest responses, followed by Evidence Yes-Bad, Evidence No-Good, and finally Evidence Yes-Good.

To compare responses, a mixed-measures ANOVA was conducted with Consequence (Bad, Good) and Evidence (No, Yes) as between-subjects factors, and Case Version (Experiment, Plant, Marble, Train) as a within-subjects factor.

Regarding the within-subjects effects, the ANOVA revealed a significant effect of Case Version ($F(1, 504) = 3.53, p = 0.02, \eta_p^2 = 0.02$). None of the interactions were significant (all p 's > 0.25). This effect size for Case Version is generally regarded as small, as one can use the same rule of thumb to interpret partial eta squared as eta squared (a small effect = 0.01, medium = 0.06, and large = 0.14; Pallant, 2010, p.

Table 1 Case Version structures

Case Version	Case
Experiment	<p>Sam is conducting research regarding 5 key claims</p> <p>Given the current direction of Sam's research, Sam will end up being [wrong or correct] about the 5 key claims</p> <p>Sam could change the direction of the research by making a particular [guess based on a coin flip or hypothesis based on past findings from reputable scientists]</p> <p>By making this [guess or hypothesis], Sam will end up being [wrong or correct] about the 5 key claims</p>
Plant	<p>Sam is categorizing 10 different plants</p> <p>Given Sam's current categorization method, Sam will end up being [wrong or correct] about the categories of all 10 plants</p> <p>Sam could change his strategy by making a [guess based on a coin flip or hypothesis based on information found in a botany book]</p> <p>By making this [guess or hypothesis], Sam will end up being [wrong or correct] about the categories of the 10 plants [by chance or because the information from the botany book was accurate or wrong]</p>
Marble	<p>Sam is estimating the number of blue marbles in a jar filled with marbles of many different colors</p> <p>Given Sam's current mathematical thinking, Sam will end up being [wrong or correct] about the number of blue marbles in the jar</p> <p>Sam could change his thinking by making a [guess based on a coin flip or hypothesis based on information from a textbook on applied mathematics]</p> <p>By making this [guess or hypothesis], Sam will end up being [wrong or correct] about the number of blue marbles [by chance or because the information from the textbook was accurate or wrong]</p>
Train	<p>Sam is recording the planned arrival times for 7 trains</p> <p>Given the train times Sam has recorded, Sam will end up being [wrong or correct] about the arrival times for all 7 trains</p> <p>Sam could modify the train times recorded in accordance with a [guess based on a coin flip or a hypothesis based on an internet search]</p> <p>By making this [guess or hypothesis], Sam will end up being [correct or wrong] about the arrival times of all 7 trains [by chance or because the information online was accurate or wrong]</p>

210).⁹ Post hoc comparisons with Bonferroni corrections revealed that the difference between participants' responses to Experiment and Train, while small, was significant

⁹ Eta squared is calculated as the amount of variance explained by the factor of interest divided by the total amount of variance. One problem with eta squared is that by adding more factors into a design one decreases the amount of variance any individual factor may account for, and in so doing may make it difficult to compare the effects of a single factor across studies. Partial eta squared overcomes this problem by using a different denominator. Partial eta squared is calculated as the amount of variance explained by the factor of interest divided by that factor's variance plus any error variance. [Levine and Hullett \(2002\)](#) caution that the partial eta squared is typically greater than the eta squared as a function of the number of factors put into the equation, i.e., the discrepancy between partial eta squared and eta squared will be larger in a five-factor than two-factor ANOVA.

Table 2 Experiment 1 descriptive statistics

Case type		N	Case Version	Justified item Mean (SD)
Evidence	Consequence			
No (Guess)	Bad	41	Experiment	16.63 (27.32)
			Plant	16.15 (23.15)
			Marble	19.15 (27.75)
			Train	12.00 (21.56)
	Good	41	Experiment	65.46 (31.69)
			Plant	61.80 (32.83)
			Marble	67.63 (27.43)
			Train	58.59 (34.46)
Yes (Hypothesis)	Bad	42	Experiment	41.97 (34.83)
			Plant	34.19 (35.13)
			Marble	36.79 (33.67)
			Train	33.64 (36.65)
	Good	47	Experiment	82.62 (22.24)
			Plant	84.70 (22.46)
			Marble	81.02 (24.89)
			Train	81.87 (20.46)

Table 3 Comparisons between Case Versions

Comparison	Absolute mean difference (SD)	Significance (Bonferroni corrected)
Experiment versus plant	2.46 (1.65)	0.82
Experiment versus marble	0.53 (1.84)	1.00
Experiment versus train*	5.20 (1.76)	0.02*
Plant versus marble	1.94 (1.78)	1.00
Plant versus train	2.74 (1.60)	0.54
Marble versus train	4.67 (1.98)	0.12

*Indicates a significant difference. Differences could range from 0 to 100

($p = 0.02$). None of the other post hoc comparisons differed significantly (all p 's > 0.12). The mean differences between participants' responses to the Case Versions appear in Table 3.

Regarding the between-subjects effects, the ANOVA revealed a significant effect of Consequence ($F(1, 168) = 147.81, p < 0.001, \eta_p^2 = 0.47$) and Evidence ($F(1, 168) = 26.87, p < 0.001, \eta_p^2 = 0.14$). The interaction was not significant ($p = 0.85$). Figure 1 displays the mean responses to each case.




Fig. 1 Experiment 1: the mean Justified response for participants’ responses to each case

4.4 Experiment 1 discussion

Experiment 1’s results reveal that there are two factors that are important to people when considering the permissibility of epistemic acts: Consequence and Evidence. Notably the effects of Consequence were stronger than Evidence. Experiment 1 involves the permissibility of making guesses and hypotheses, but the philosophical literature mostly concerns the epistemic permissibility of forming beliefs or accepting propositions (Dunn 2016). Hence, to more directly engage with the literature, Experiment 2 was performed.

5 Experiment 2

Experiment 2 was performed to investigate people’s intuitions about the epistemic permissibility of forming beliefs. For each case, participants were asked whether the belief is justified. If the results of this study differ from Experiment 1, this suggests that norms concerning belief take a different form from those concerning guesses and hypotheses.

5.1 Participants

231 new participants recruited via Amazon Mechanical Turk completed the survey. We required participants to have at least a 95% HIT approval rate and at least a US high school diploma. The experiment was run until at least 40 participants in each group answered all the comprehension items correctly. Of the 231 participants, 47 were removed for answering one or more comprehension items incorrectly, and 9 were removed for having any formal academic training in philosophy whatsoever. Of the remaining 175 participants, 89 identified as female and the median age was 39.

5.2 Materials and procedure

The materials and procedure for Experiment 2 were similar to Experiment 1 except for the following details that were altered for the cases and target item.

The cases were designed to be similar to Experiment 1, with the relevant changes made as needed to produce belief cases. For example, while the Experiment 1 Case Versions formerly read, “By making this [guess or hypothesis]...” for Experiment 2 the Case Versions read, “By forming this belief...” (full versions are found in Appendix A, Table 5).

Adjustments were also made for the target item. While in Experiment 1 the target items read, “The [guess or hypothesis] is justified,” in Experiment 2 the target item read, “The belief is justified.”

5.3 Results

Means and standard deviations for all cases are displayed in Table 4. Descriptively, the Evidence No-Bad cases received the lowest ratings, followed by Evidence Yes-Bad, Evidence No-Good (though note that Evidence Yes-Bad and Evidence No-Good were very similar), and finally Evidence Yes-Good.

To compare responses, a mixed-measures ANOVA was conducted with Consequence (Bad, Good) and Evidence (No, Yes) as between-subjects factors, and Case Version (Experiment, Plant, Marble, Train) as a within-subjects factor.

Regarding the within-subjects effects, the ANOVA revealed a significant three-way interaction ($F(1, 513) = 4.83, p = 0.03, \eta_p^2 = 0.03$). No other significant within-subjects effects were found (all p 's > 0.06). The interaction was further examined graphically (see Fig. 2). Looking at the Evidence No conditions, one sees that the participants' judgements were always lower when the consequences were bad than when they were good, and that Case Version had little effect on this relationship. Likewise, looking at the Evidence Yes conditions, one again sees that the participants' judgements were always lower when the consequences were bad than when they were good, but dissimilarly the magnitude of this difference depended on Case Version. Specifically, the largest difference appears between Bad and Good for Train and the smallest difference appears for Experiment. As the three-way interaction is small and is not especially relevant to this paper, it will not be discussed further.

Regarding the between-subjects effects, the ANOVA revealed a significant effect of Consequence ($F(1, 171) = 64.41, p < 0.001, \eta_p^2 = 0.27$), Evidence ($F(1, 171) =$

Table 4 Experiment 2 descriptive statistics

Case type		N	Case Version	Justified item Mean (SD)
Evidence	Consequence			
No	Bad	43	Experiment	25.95 (27.54)
			Plant	26.72 (29.15)
			Marble	26.93 (28.69)
			Train	26.00 (28.92)
	Good	45	Experiment	43.62 (29.18)
			Plant	41.29 (28.86)
			Marble	42.47 (29.10)
			Train	39.42 (29.47)
Yes	Bad	47	Experiment	48.28 (29.93)
			Plant	38.43 (35.13)
			Marble	43.60 (33.67)
			Train	34.58 (36.65)
	Good	40	Experiment	82.31 (22.93)
			Plant	88.81 (17.85)
			Marble	82.21 (21.07)
			Train	83.93 (20.89)

62.3, $p < 0.001$, $\eta_p^2 = 0.27$), and a significant interaction between Consequence and Evidence ($F(1, 171) = 14.60$, $p < 0.001$, $\eta_p^2 = 0.08$). The interaction was such that the difference between participants' responses to Consequence (Bad vs. Good) was greater in the Evidence Yes conditions. Figure 2 displays the mean responses to each case.

5.4 Experiment 2 discussion

Experiment 2's results reveal that, like Experiment 1, there are two factors that are important to people when considering epistemic permissibility: consequence and evidence. Unlike Experiment 1, when considering the permissibility of beliefs (rather than guesses or hypotheses), the effect of Consequence was as strong as the effect of Evidence. Also unlike Experiment 1, an interaction emerged between the factors Consequence and Evidence.

6 General discussion

This discussion section has four main parts. First, we provide our preferred explanation of our results. Second, we consider how our results bear on those of Andow's previous study. Third, we articulate some implications of our results for philosophy. Fourth, we look at some speculative hypotheses based on our current results and avenues for




Fig. 2 Experiment 2: the mean Justified response for participants' responses to each case

further research. Finally, we take some time to anticipate and respond to a number of objections.

6.1 Preferred explanation

Grant for the moment that our target item (Justified) measures participants' intuitions about epistemic normativity (but see "Objections and replies"). If so, our results for both Experiment 1 and 2 suggest that lay people do not have intuitions that accord with a view like DEC: direct consequentialists should not treat the two good consequence conditions in such different ways, for the stated consequences of the two relevant conditions do not differ in any regard. On the other hand, neither do our results suggest that lay intuitions accord with epistemic deontology: deontologists should not treat the two conditions where there is evidence (the Evidence Yes conditions) in such different ways; these conditions only differ in terms of consequences, and deontologists should not consider such differences to be relevant.

What our results suggest instead is that both the maximization of the epistemic good and the presence/absence of evidence are important to lay people when considering whether an act is epistemically justified. Thus, we propose a model of lay intuitions according to which there are two factors that underlie lay intuitions about epistemic

permissibility: (1) the extent to which an epistemic act results in the maximization of the epistemic good (e.g., true beliefs), and (2) the extent to which an epistemic act (e.g., guessing, hypothesizing, believing) enjoys independent support (e.g., evidence). Interestingly, our results suggest that the weight that these two factors are given depends on the nature of the epistemic act under consideration. For guesses and hypotheses, maximization of the epistemic good is given more weight than evidence, whereas for beliefs neither factor is dominant.

6.2 Relation to Andow's study

Andow's (2016a) results might seem to conflict with ours. All the cases used in Andow's paper fall into one of two categories: cases in which a researcher can change some equipment or cases in which a researcher can make an assumption (in which it is never mentioned that there is evidence and sometimes even stated that there is not). Across Andow's two experiments, he develops five cases that fit the latter category (assumption cases) and in all such cases participants consider the epistemic act of making an assumption to be impermissible (means between 25 and 41). Our Evidence No-Good cases were designed to be roughly analogous to Andow's assumption cases, but our participants' responses were generally higher (means between 39 and 68).

Although Andow's results do seem to conflict with ours, we caution against drawing any strong conclusions from the descriptive differences. We are inclined to think that Andow's findings do not present much of challenge to the two-factor model of lay intuitions that we proposed (more relevant to this in "Speculative hypotheses and future research"). There are many reasons why Andow's results might be different descriptively. The differences could be due to the studies having different designs, or to the studies making use of different participant pools (Prolific vs. Amazon Turk participants). One notable, potential explanation is that the cases involved in our experiments and in Andow's describe different epistemic acts. Our results reflect intuitions about guesses, hypotheses, and beliefs, whereas Andow's results reflect intuitions about assumptions and changing equipment (though the latter is not relevant here). Thus, one possible explanation of the differences might be that lay people treat assumptions as being particularly epistemically dubious.

6.3 Philosophical implications

Let us assume for the sake of discussion that intuitions are indeed relevant to which epistemic theory is true. Let us also assume that lay intuitions are relevant to deciding whether a philosophical view is intuitive or not, at least in the relevant areas. We can now revisit the intuition based objection presented earlier, in Sect. 2, which went roughly like this:

A form of epistemic consequentialism that endorses a deontic theory like DEC is committed to the idea that it is epistemically required to accept the Truth Fairy's offer (supposing a veritist account of final value). Intuitively, however, it is not *epistemically* acceptable to believe something that one has no independent

reason to believe, simply because of the consequences. Thus, such forms of epistemic consequentialism are unintuitive and so should be rejected.

Our results add weight to the idea that the defender of a view like DEC has a heavy counterintuitive burden to bear and so supports the above objection. Our evidence is stronger than Andow's, as we looked at intuitions concerning belief, in addition to guesses and hypotheses. Intriguingly, our results suggest that a straightforward deontological account has a heavy counterintuitive burden too. The results of Experiment 2 suggest that the maximization of an individual's epistemic position and independent motivation are of equal importance to belief justification. Indeed, it is only when a belief will maximize an individual's epistemic position in addition to having independent support that our participants strongly agreed that it is justified. The results of Experiment 1, concerning guesses and hypotheses, suggest that the maximization of an individual's epistemic position is more important than independent motivation but that both are relevant to justification.

6.4 Speculative hypotheses and future research

While our results do not speak in favor of any one specific deontic theory, we can nonetheless say more about the overall picture of lay epistemic permissibility that our results suggest. The picture is compatible with both monistic and pluralistic accounts of final value. The lay deontic picture best supported by our results is one whereby two things hold:

1. There is a two-factor structure for all types of epistemic act: one factor concerns maximization and the other concerns independent motivation.
2. Slightly different deontic theories apply for different types of epistemic act.

Our results support (1), because both maximization and independent motivation were important to our participants when deciding whether an epistemic act is justified. Our results support (2), because the relative importance of consequences and independent motivation depended on the epistemic act we asked participants to consider.

The two-factor structure that we consider could be due to a dyadic account of final value or to a dyadic deontic theory. It is possible that lay people hold a pluralistic view of final value and treat both (i) true beliefs (and perhaps guesses, hypotheses, etc.) and (ii) independent support as being of final value. These values might be regarded as incommensurable such that an action that results in more, for example, true beliefs has greater overall value in one respect (consequentially), and an action involving independent support has greater overall value in another (evidentially). Such a pluralistic account of final value would presumably be accompanied by a deontic theory that took both values into account (e.g., maximize both or maximize at least one). However, perhaps instead, the lay conception of final value might recognize only true beliefs; independent support would be only of instrumental epistemic value. Such a monistic account of final value is consistent with a two-factor model: the two factors can enter the picture at the level of the deontic theory (we will look at some examples shortly).

We propose the following speculative hypothesis for how (1) and (2), as defined above, apply specifically for guesses, hypotheses, and beliefs:

- Guesses: a guess can be justified despite not enjoying any independent support so long as it maximizes.
- Hypotheses: having independent support is not sufficient for a hypothesis to be justified (although it may have a positive influence on levels of justification); the hypothesis also needs to maximize.
- Beliefs: a belief must maximize and enjoy independent support to be justified. Both factors independently have a positive influence on levels of justification.

We might broaden this speculative hypothesis using Andow's (2016a) results with the addition of a slightly different deontic theory for assumptions as follows:

- Assumptions: maximizing is not sufficient for an assumption to be justified; it also needs to have independent support.

There are many ways to go about testing this speculative hypothesis. Which one a researcher should pursue depends on many factors, for example practicality, including monetary expense. This said, one method would be to run a large experiment with a similar structure to Experiments 1 and 2 but with the addition of the type of epistemic act (guesses, hypotheses, beliefs, and assumptions) as a third factor. This would produce a $4 \times 2 \times 2$ design. Such an experiment could use similar cases to those we used, but it could also use new cases to explore the generalizability of the patterns of results we observed (the less alterations from our current experiments the more able to engage with our results). The objections and responses to the current experiments, which we address in "Objections and replies," should be kept in mind when designing future experiments.

We will now take some time to consider a more specific issue: what our results tell us about the lay understanding regarding the normativity of belief. When discussing epistemic normativity, philosophers usually focus on the permissibility of beliefs. While we are inclined to think that a more general approach, which looks at a greater variety of epistemic actions, is more valuable, we will clarify the take home message of our results for those who are primarily interested in beliefs. Again, our current results do not speak in favor of any one specific account of the lay understanding of epistemic normativity (hereon "lay accounts"). However, they do help to narrow the field. In the rest of this section, we will first outline three different, potentially true lay accounts compatible with our results. We will then consider how one might go about empirically deciding between them.

The first lay account is a pluralistic one with two necessary conditions:

Pluralistic, with two necessary conditions (P2NC): a belief that p is epistemically permissible for a subject only if (i) believing p maximizes epistemic value for the subject, and (ii) the subject's belief that p enjoys independent support.

P2NC represents an important alternative to simple direct consequentialist and deontological lay accounts, in placing importance on both consequences and independent motivation.

The second lay account places less importance on the distinction between permissibility and impermissibility using necessary conditions, but has a certain isomorphism with P2NC. It is a pluralistic lay account with two multiplicative factors:

Pluralistic, with two multiplicative factors (P2MF): A belief that p is more epistemically permissible for a subject if it has (i) better epistemic consequences, or (ii) enjoys independent support, and (iii) the extent to which having better epistemic consequences leads to greater permissibility is greater the more independent support.

The third lay account tries to maintain a purely consequentialist and purely veritist explanation of epistemic value. It is thus a monistic lay account, and has two necessary conditions (one direct consequentialist condition and one indirect):

Monistic, with two necessary conditions (M2NC): A belief that p is epistemically permissible for a subject only if (i) believing p maximizes true beliefs for the subject, and (ii) the subject's believing that p conforms to a rule for belief formation, the adoption of which would maximize true beliefs for the subject, where "only form beliefs that have independent support" is such a maximizing rule.

How might future empirical research decide between these lay accounts? The three lay accounts differ along two dimensions: (a) some are monistic about final epistemic value, while others are pluralistic; (b) some exploit necessary conditions for permissibility, while others involve linear relationships between permissibility and the factors which influence permissibility.¹⁰ How might one investigate whether the folk entertain monist or pluralist views? One would need to distinguish between participants who are sensitive to a particular factor (e.g., evidence) and who treat that factor as having final epistemic value, and participants who are sensitive to that same factor but treat it as having only instrumental epistemic value. We think that distinguishing between such participants without qualitative insights will be very difficult. Thus, to determine whether the folk are monists or pluralists, we propose the use of focus group discussions and/or interviews or similar.¹¹

How might one investigate whether the folk understand the normativity of belief to employ necessary conditions or a more linear relationship? What follows shows how one can, in principle, empirically determine whether the lay account employs a necessary condition or a more linear relationship with respect to a single variable.^{12, 13}

¹⁰ There is a fourth logically coherent account: monist with sensitivity to two multiplicative factors. For instance, one of the two factors may be seen as being of only non-final value.

¹¹ See [Andow \(2016c\)](#) for more details about how qualitative methods can be used in experimental philosophy.

¹² For an example of a study using a similar method to probe intuitions about moral acceptability, see [Shenhav and Greene \(2010\)](#).

¹³ For simplicity, we focus on an example involving a single independent variable. Note that in the lay accounts that matter for the current paper, there are two independent variables of interest, because the lay accounts involve a dyadic structure. This would make the matter of collecting the relevant data and fitting nonlinear models to the data more complicated. However, we only wish to illustrate how the two types of lay account are in principle empirically distinguishable, and we encourage other researchers to look into this.




Fig. 3 An illustration of linear and sigmoidal regression lines for a fictional data set

We propose collecting responses from participants at a range of values of this variable. We propose then looking at whether the relationship between the dependent and independent variable is best captured by a linear function or a nonlinear function (e.g., modelled using sigmoidal curves). One should determine whether a linear regression model or a nonlinear regression model (e.g., one that uses a sigmoidal function) better predicts the data. If a nonlinear sigmoidal model (or similar) were a better predictor, this would be a sign that the participants were operating with a conception that employs thresholds, e.g., necessary conditions. See Fig. 3 for an illustration of a hypothetical set of findings depicting nonlinear sigmoidal and linear regression lines.

6.5 Objections and replies

Before closing, let us say something to anticipate and respond to a few possible worries about our results and our interpretation of them. We will look at five objections.

Objection 1: these results likely reflect some moral attitudes rather than participants' intuitions about epistemic normativity. This is a genuine concern about empirical research in this domain. Andow (2016a) discusses various reasons why it is difficult to design questions that clearly probe intuitions about epistemic normativity rather than, for example, intuitions about what is morally permissible. The probes we use employ the word “justified,” which frequently occurs in moral discourse. So, it may be that participants are reporting their moral intuitions rather than their epistemic intuitions. However, we are inclined to think that this is not the case. If our participants are reporting their moral intuitions, then, for example, they must have strong moral objections to guessing on the basis of no evidence when doing so will not maximize one’s epistemic position, but have no such moral objections when doing so will maximize one’s epistemic position. This strikes us as rather implausible. The issues at stake and the overall context seem clearly epistemic and not moral. This said, this seems to be an empirically tractable issue, and we welcome attempts to investigate it further.

*Objection 2: for guesses and hypotheses in particular it is plausible that participants assessed pragmatic rather than epistemic justification.*¹⁴ All of the cases we used were purposefully written to depict scenarios in which there is little practically at stake for the subject. For example, at the train station, Sam is not trying to catch a train but is simply recording train arrival times. The cases do not explicitly exclude the possibility that there is much practically at stake, but they do not invite any such interpretation either. So, we are not too concerned that our results track participants' intuitions about pragmatic justification. Nonetheless, we can point to nothing about our results that would definitively rule out such a possibility. Intuitions about both epistemic and pragmatic justification might be influenced by both considerations of epistemic utility and of independent support. We welcome further research on this issue. We suspect that some form of protocol analysis may be needed to firmly ascertain whether participants are assessing the permissibility of acts epistemically or pragmatically, or indeed whether they clearly distinguish the two.¹⁵

Objection 3: the study is concerned specifically with the epistemic acceptability of particular acts (e.g., forming a belief), but the statements that participants were ultimately asked about might have been interpreted as being about epistemic objects (the beliefs themselves). The reason for this concern is that the statements literally asked about "the guess," "the hypothesis," and "the belief," not "the subject's making the guess," "the subject's making the hypothesis," and "the subject's forming the belief." We did this because we felt it was more vernacular and easier to read and understand. However, what we did could be seen as a problem, because perhaps in some objective sense, the belief, for instance, is justified, even if the particular subject is not justified in forming it. We appreciate this concern, but we think it unlikely that our participants were confused about what we meant. They had just read a case about making a guess or hypothesis or forming a belief, and they had to answer comprehension questions about the case. The relevant case continued to be displayed when they arrived at the page asking for their response to the target item (Justified).

Objection 4: Case Type was a between-subjects variable, and so we cannot comment on whether there are order effects of Case Type. This is true and a limitation of our study. The study was designed this way deliberately. The reason was to avoid the complication of order effects for Case Type before we have understood the basic components underlying participants' intuitions about epistemic normativity. Experimental research using ethical trolley cases, which are analogous in notable respects

¹⁴ A related objection is that epistemic notions of justification do not even apply to guesses in ordinary discourse. This sounds like an empirical question. However, in the opinion of the native English speaking authors of this article, it makes perfect sense to talk about a good or bad guess in an epistemic sense. "Jeff lost the Poker game because of a bad guess based on miscounting the number of kings played." Here the rationale for why the guess is bad is epistemic not pragmatic. "He won because of a good guess based on his impeccable knowledge of the game." Here the rationale for the guess being good is epistemic.

¹⁵ Colman defines protocol analysis thus: "Any of a number of techniques of qualitative research in which thought processes elicited from research participants or subjects while they were performing a mental task are studied and classified, often by means of content analysis" (2016).

to our cases, has revealed robust order effects (Wiegmann et al. 2012). It is an open question whether intuitions about epistemic normativity exhibit similar effects. Having investigated the basic components underlying participants' intuitions, we think a reasonable course of action for future research is to examine order effects. It would be particularly intriguing to see whether they are similar to the order effects known for moral cases.¹⁶

Objection 5: what is and is not intuitive to non-philosophers is of no relevance to philosophical epistemology. At least one of the current authors is inclined to think that this objection is deeply mistaken: a central part of philosophy is and should be an attempt to make sense of our ordinary pre-theoretic ways of thinking about philosophically interesting phenomena such as epistemic normativity (Andow 2016b). Given this stance, it should be clear why one would think a proper empirically-informed understanding of how non-philosophers think about epistemic normativity is important to philosophical epistemology. Another of the current authors does not think that a *central part* of philosophy is (or should be) to make sense of our ordinary pre-theoretic ways of thinking but still disagrees with the objection, in that he thinks making sense of our pre-theoretic ways of thinking is an instrumental part of philosophy. Regardless, such metaphilosophical worries are outside the scope of this article. Those concerned with such worries must recognize that many epistemologists think that theorizing in epistemology should be sensitive to intuitions.¹⁷ Hence, our results provide an important source of information for philosophers working in this area.

Acknowledgements We would like to thank Ivar Hannikainen, Aimie Hope, and Luke Elson for helpful comments. We are also grateful to two fantastic anonymous referees.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

¹⁶ Experimenters who want to investigate order effects for Case Type may consider running participants in all the orders, and then comparing their responses across orders. Before doing this, note that the number of orders increases as one considers more Case Types. With just two (A, B) there are two orders (e.g., A then B, and B then A), and with four Case Types there are 24 orders. Recruiting enough participants to run and compare all the orders may prove untenable. To overcome this, one may use a balanced Latin squares design (Richardson et al. 2011). A balanced Latin squares design limits the number of orders considered in an experiment so that each Case Type appears only once in each position in the ordering sequence and that each Case Type follows every other Case Type an equal number of times.

¹⁷ See Nagel (2007) for discussion of various positions on this point pertaining to epistemology more generally. See Elstein and Jenkins (forthcoming) for a discussion of epistemic consequentialism that is conducted very explicitly in terms of intuitions.

7 Appendix A

Table 5 Case Version structures for belief

Case Version	Case
Experiment	<p>Sam is conducting research regarding 5 key claims</p> <p>Given the current direction of Sam's research, Sam will end up being [wrong or correct] about the 5 key claims</p> <p>Sam could change the direction of the research by forming a particular belief based on [a coin flip or past findings from reputable scientists]</p> <p>By forming this belief, Sam will end up being [wrong or correct] about the 5 key claims</p>
Plant	<p>Sam is categorizing 10 different plants</p> <p>Given Sam's current categorization method, Sam will end up being [wrong or correct] about the categories of all 10 plants</p> <p>Sam could change his strategy by forming a belief based on [a coin flip or information found in a botany book]</p> <p>By forming this belief, Sam will end up being [wrong or correct] about the categories of the 10 plants [by chance or because the information from the botany book was accurate or wrong]</p>
Marble	<p>Sam is categorizing 10 different plants</p> <p>Given Sam's current categorization method, Sam will end up being [wrong or correct] about the categories of all 10 plants</p> <p>Sam could change his strategy by forming a belief based on [a coin flip or information from a textbook on applied mathematics]</p> <p>By forming this belief, Sam will end up being [wrong or correct] about the categories of the 10 plants by chance. [by chance or because the information from the textbook was accurate or wrong]</p>
Train	<p>Sam is recording the planned arrival times for 7 trains</p> <p>Given the train times Sam has recorded, Sam will end up being [wrong or correct] about the arrival times for all 7 trains</p> <p>Sam could modify the train times recorded by forming a belief based on [a coin flip or an internet search]</p> <p>By forming this belief, Sam will end up being [wrong or correct] about the arrival times of all 7 trains [by chance or because the information online was accurate]</p>

References

- Ahlstrom-Vij, K., & Dunn, J. (2014). A defence of epistemic consequentialism. *The Philosophical Quarterly*, 64, 541–551.
- Ahlstrom-Vij, K., & Dunn, J. (forthcoming). *Epistemic consequentialism*. Oxford: Oxford University Press.
- Alexander, J., & Weinberg, J. (2007). Analytic epistemology and experimental philosophy. *Philosophy Compass*, 2, 56–80.
- Andow, J. (2016a). Do non-philosophers think epistemic consequentialism is counterintuitive? *Synthese*. doi:10.1007/s11229-016-1071-7.
- Andow, J. (2016b). Intuitions. *Analysis*, 76, 232–246. doi:10.1093/analys/anv062.
- Andow, J. (2016c). Qualitative tools and experimental philosophy. *Philosophical Psychology*, 29, 1128–1141. doi:10.1080/09515089.2016.1224826.

- Andow, J. (forthcoming). Epistemic consequentialism, truth fairies and worse fairies. *Philosophia*. Accepted March 6, 2017.
- Berker, S. (2013). The rejection of epistemic consequentialism. *Philosophical Issues*, 23, 363–387.
- Colman, A. M. (Ed.). (2016). Protocol analysis. In *A dictionary of psychology* (4th ed.). Oxford: Oxford University Press.
- Dunn, J. (2016). Epistemic consequentialism. In *The internet encyclopedia of philosophy*. <http://www.iep.utm.edu/>. Accessed June 21, 2016.
- Dunn, J., & Ahlstrom-Vij, K. (forthcoming). Is reliabilism a form of consequentialism? *American Philosophical Quarterly*. ISSN 0003-0481. E-ISSN 2152-1123.
- Elstein, D. & Jenkins, C. I. (Forthcoming). The Truth Fairy and the Indirect Epistemic Consequentialist. In N. J. L. L. Pedersen & P. Graham (Eds.), *Epistemic entitlement*. Oxford: Oxford University Press.
- Firth, R. (1978). The Schneck lectures. Delivered at the University of Arizona in 1978. First published in Firth, 1998, pp. 315–367.
- Firth, R. (1981). Epistemic merit, intrinsic and instrumental. *Proceedings and Addresses of the American Philosophical Association*, 55, 5–23.
- Foot, P. (1967). The problem of abortion and the doctrine of double effect. *Oxford Review*, 5, 5–15.
- Goldman, A. (2007). Philosophical intuitions: Their target, their source, and their epistemic status. *Grazer Philosophische Studien*, 74, 1–26.
- Goldman, A. (2015). Reliabilism, veritism, and epistemic consequentialism. *Episteme*, 12, 131–143.
- Greene, J. (2016). Solving the trolley problem. In J. Sytsma & W. Buckwalter (Eds.), *A companion to experimental philosophy* (pp. 175–178). West Sussex: Wiley.
- Hedberg, T. (2014). Epistemic supererogation and its implications. *Synthese*, 191, 3621–3637.
- Jenkins, C. I. (2007). Entitlement and rationality. *Synthese*, 157, 25–45.
- Klausen, S. H. (2009). Two notions of epistemic normativity. *Theoria*, 75, 161–178.
- Levine, T. R., & Hullett, C. R. (2002). Eta squared, partial eta squared and the misreporting of effect size in communication research. *Human Communication Research*, 28, 612–625.
- Nado, J. (2014). Philosophical expertise. *Philosophy. Compass*, 9, 631–641.
- Nagel, J. (2007). Epistemic intuitions. *Philosophy. Compass*, 2, 792–819.
- Pallant, J. (2010). *SPSS Survival Manual: A step by step guide to data analysis using SPSS* (4th ed.). Berkshire, England: Open University Press.
- Percival, P. (2002). Epistemic consequentialism-I. In *Proceedings of the Aristotelian Society, Supplementary Volumes* (Vol. 76, pp. 121–151).
- Pritchard, D. (2014). Sceptical intuitions. In A. R. Booth & D. P. Rowbottom (Eds.), *Intuitions*. Oxford: Oxford University Press.
- Richardson, P., Goodwin, A., & Vine, E. (2011). *Research methods and design in psychology*. Exeter: Shortun Press Ltd.
- Shenhav, A., & Greene, J. (2010). Moral judgments recruit domain-general valuation mechanisms to integrate representations of probability and magnitude. *Neuron*, 67, 667–677.
- Thomson, J. J. (1971). A defense of abortion. *Philosophy and Public Affairs*, 1, 47–66.
- Thomson, J. J. (1976). Killing, letting die, and the trolley problem. *The Monist*, 59, 204–217.
- Wiegmann, A., Okan, Y., & Nagel, J. (2012). Order effects in moral judgment. *Philosophical Psychology*, 25, 813–836.