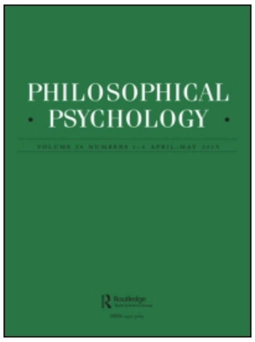

Routledge
Taylor & Francis Group

Philosophical Psychology

ISSN: 0951-5089 (Print) 1465-394X (Online) Journal homepage: http://www.tandfonline.com/loi/cphp20

Lottery judgments: A philosophical and experimental study

Philip A. Ebert, Martin Smith & Ian Durbach

To cite this article: Philip A. Ebert, Martin Smith & Ian Durbach (2018) Lottery judgments: A philosophical and experimental study, Philosophical Psychology, 31:1, 110-138, DOI: 10.1080/09515089.2017.1367767

To link to this article: https://doi.org/10.1080/09515089.2017.1367767

Published online: 21 Sep 2017.

PHILOSOPHICAL PSYCHOLOGY, 2018
VOL. 31, NO. 1, 110–138
https://doi.org/10.1080/09515089.2017.1367767

Lottery judgments: A philosophical and experimental study

Philip A. Ebert[a], Martin Smith[b] and Ian Durbach[c,d]

[a]Department of Law and Philosophy, University of Stirling, Stirling, UK; [b]University of Edinburgh, Edinburgh, UK; [c]Department of Statistical Sciences, University of Cape Town, Cape Town, South Africa; [d]African Institute for Mathematical Sciences, Cape Town, South Africa

ABSTRACT

In this paper, we present the results of two surveys that investigate subjects' judgments about what can be known or justifiably believed about lottery outcomes on the basis of statistical evidence, testimonial evidence, and "mixed" evidence, while considering possible anchoring and priming effects. We discuss these results in light of seven distinct hypotheses that capture various claims made by philosophers about lay people's lottery judgments. We conclude by summarizing the main findings, pointing to future research, and comparing our findings to recent studies by Turri and Friedman.

ARTICLE HISTORY

Received 8 February 2017
Accepted 12 June 2017

KEYWORDS

Behavioral decision-making; justified belief; knowledge; lottery proposition; statistical evidence; testimonial evidence

1. Introduction

Consider the proposition that a particular ticket has lost a lottery. Call this a *lottery proposition*. Many philosophers report the intuition that one cannot know a lottery proposition based on purely statistical evidence – one cannot know that a ticket has lost based purely on the overwhelming odds against it winning. This "lottery intuition" has exerted considerable influence in recent epistemology. It has been endorsed by Cohen (1998), Dretske (1970), Nelkin (2000), Nozick (1981), Pritchard (2005), Smith (2010), Vogel (1990), Williamson (2000), Williamson (2009b), and many others.[1] Those who endorse the lottery intuition often take it to be widely shared among philosophers and "lay people" alike. So, for example, John Hawthorne writes in the introduction to *Knowledge and Lotteries*:

> And yet I take it as a datum that there is a strong inclination to claim that the relevant lottery propositions are not known. Nor is this merely a datum about the inclinations of philosophers. After all, the motto of the New York State lottery is 'Hey, you never know'. (Hawthorne, 2004, p.8)

CONTACT Philip A. Ebert p.a.ebert@stir.ac.uk

In a similar fashion, Peter Baumann writes:

> Almost everybody seems to agree that nobody can know the outcome of a lottery in advance or just on the basis of the statistical evidence (even if that evidence is overwhelmingly in favor of a certain outcome). (Baumann, 2004, p. 416)

The lottery intuition plays a wide-ranging role in current epistemological theorizing. So, for example, it is often appealed to as an adequacy constraint to motivate or test various conditions on knowledge, such as safety (anti-luck), sensitivity, reliabilist, or normalcy conditions. Also, the lottery intuition features in a number of well-known epistemic puzzles[2] which epistemologists typically try to solve in such a way as to leave the lottery intuition intact. Moreover, it has played an important role in the debate over norms of assertion, and figures centrally in Williamson's widely discussed defense of the *knowledge norm*.[3]

As well as denying that one can know a lottery proposition based purely on the odds against it winning, some philosophers have gone further and denied that one can even *justifiably believe* a lottery proposition on this basis (Nelkin, 2000; Sutton, 2007; more recently, Smith, 2010, 2016; Smithies, 2012). Those who endorse this claim, however, rarely describe it as "intuitive" or widely accepted, and almost universally regard it as something standing in need of substantial argument.

Lay judgments about what we can know and justifiably believe about lottery outcomes have recently been scrutinized in two experimental philosophy studies by Turri and Friedman (2014) and Friedman and Turri (2015). Turri and Friedman found that an impressive majority of subjects (~90%) do share the lottery intuition, judging that one cannot know a lottery proposition based purely on statistical evidence. The studies also showed that a nearly as impressive majority of subjects (~80%) judged that one could justifiably believe a lottery proposition purely on the basis of statistical evidence. Turri and Friedman take these findings to weigh against certain views about lottery judgments, such as what they call the "justification account." According to this account, the reason people judge that one cannot know that a lottery ticket is a loser based purely on statistical evidence is because they judge that one cannot justifiably believe that a lottery ticket is a loser based purely on statistical evidence, and take justification to be necessary for knowledge. As Turri and Friedman claim:

> The justification account says that in basic lottery cases people deny knowledge because they think justification is absent. So if the justification account is correct, very few participants should say that Lois is justified in thinking that the ticket is a loser … The results were highly unfavorable to the justification account. (Turri & Friedman, 2014, pp. 49–50)

We agree that Turri and Friedman's findings weigh against the justification account, but are very doubtful whether the account has ever been endorsed in quite the way that Turri and Friedman describe.[4]

Turri and Friedman also suggest that their findings may put pressure on certain substantial views about the nature of knowledge and justification and the relationship between the two. They single out in particular the recently prominent "knowledge account" of justification, on which knowledge and justified belief are taken to be identical:

> Our results … should be taken into consideration when evaluating the increasingly popular "knowledge-first" approach in epistemology … Perhaps the most radical plank in the knowledge-first platform is the identification of justification with knowledge (Sutton, 2007). To the extent that this is supposed to reflect the way people actually think about knowledge, our results undermine the view. (Turri & Friedman, 2014, p. 32)

Again, we are inclined to doubt that Turri and Friedman's results, or indeed our results, directly bear against the knowledge account, mainly because it is doubtful that its proponents consider the view to be genuinely descriptive of how lay people judge the epistemic status of lottery propositions. Sutton, in fact, acknowledges that his view clashes with ordinary judgments about when lottery propositions can be justifiably believed, and he offers an explanation for why people judge the way they do – we will discuss his proposal in more detail below.

Now, of course, there is a general debate about the extent to which lay intuitions about philosophical thought experiments bear upon philosophical theorizing of this sort. Some critics suggest that lay intuitions are less trustworthy than those of experts – that is, professional philosophers.[5] Others have taken an even harder line and denied that intuitions have any general relevance for the evaluation of philosophical claims – be they the intuitions of lay people or professional philosophers.[6] Here we won't take a stance on this issue, nor do we have to take a stance, since our investigation has a different focus. Philosophers writing in epistemology have not shied away from claims that are explicitly *about* lay people's judgments and which can be straightforwardly tested by surveying these judgments. We have seen examples of such claims in the above quotes, and more will be discussed below. One might deny that these are proper "philosophical" claims (their appearance in philosophical works notwithstanding) but we won't attempt to enforce any such boundary here.[7]

The main aim of this paper, then, is to investigate lay people's judgments about lotteries. As well as attempting to replicate the results of previous studies, we extended these studies in various ways. Drawing upon recent psychological studies, we also tested whether lottery judgments are sensitive to the size of the lottery involved and to the strength of the statistical evidence presented. Relatedly, we undertook to test whether lottery judgments exhibited an anchoring effect when subjects were asked about a sequence of lotteries either increasing or decreasing in size. And, finally, by drawing on recent philosophical discussion of the lottery cases, we tested for whether lottery judgments were subject to a priming effect in light of a distinction between the proposition that a given ticket is a loser and the proposition that a given ticket is *probably* a loser.

In order to structure the discussion of our methodology and our results, we will first identify and provide *prima facie* motivation for seven distinct hypotheses capturing numerous claims made by philosophers concerning lay people's judgments about lotteries. In the following section, we explain how we tested for these hypotheses by outlining the details of our surveys and offer some motivation for our preferred statistical methodology. In Section 4, we present and discuss the main findings from our surveys, followed by a discussion of these results in relation to the seven hypotheses. We close by comparing our results to the other two studies, assess potential repercussions of our findings, and suggest some future directions for further studies on this topic.

2. Motivating seven hypotheses

We begin with two hypotheses strongly supported by Turri and Friedman's results:

> ***Hypothesis 1:*** *No-knowledge*
> People have a general tendency to deny knowledge of lottery propositions when the evidence is statistical.

> ***Hypothesis 2:*** *Yes-Justification*
> People have a general tendency to ascribe justified belief in lottery propositions when the evidence is statistical.

We assume a fair and decent sized lottery (more on this below when we outline our vignettes). Importantly, the assumption here is that the belief in the lottery proposition is supported *only* by the relevant statistical evidence derived from the setup of the lottery.

The third hypothesis can be motivated by appeal to a well-known study in psychology, Wells (1992), which has established the aptly named "Wells effect": subjects are very reluctant to make liability decisions based on "naked" statistics, that is, purely statistical evidence, in a simulated court case. In contrast, subjects are often willing to make liability judgments based on testimonial evidence, even if, from a probabilistic perspective, the relevant testimony is less probative than the statistical evidence. This leads us to our third hypothesis to be tested (which was also tested in detail in Turri & Friedman, 2014; Friedman & Turri, 2015).

> ***Hypothesis 3:*** *Testimonial*
> People have a stronger tendency to ascribe knowledge of lottery propositions when the evidence is testimonial as opposed to when it is statistical.

While the previous two studies concerned propositions expressed by sentences of the form "this ticket is a losing ticket" – propositions that we label *lottery propositions* (LP) – it is important to also consider what we call *probabilistic lottery propositions* (PLP), expressed by sentences of the form: 'it is very likely that this ticket is a losing ticket.' Many of the philosophers who have denied that statistical evidence can provide justification for believing lottery propositions grant that it can nevertheless provide justification for believing probabilistic lottery propositions – that one can justifiably believe that a ticket has very likely lost a lottery based only on statistical evidence (compare Nelkin, 2000; Smith, 2010, 2016; Smithies, 2012; Sutton, 2007).

Hypothesis 4: *Probabilistic*
People have a stronger tendency to ascribe justified belief in probabilistic lottery propositions as opposed to categorical lottery propositions, and this effect will be particularly pronounced when the evidence is statistical.

As noted above, philosophers who defend the claim that one cannot justifiably believe lottery propositions on the basis of statistical evidence don't generally take the claim to be "intuitive" and often concede that many people have the *opposite* intuition – that one *can* justifiably believe lottery propositions on the basis of statistical evidence. Some of these philosophers, however, write as though this contrary intuition may weaken once people are made aware of the distinction between lottery propositions and probabilistic lottery propositions (Nelkin, 2000; Smith, 2010, 2016; Smithies, 2012). Sutton (2007), however, is more explicit, suggesting that a failure to distinguish between lottery propositions and probabilistic lottery propositions may be directly responsible for philosophers and lay people's tendency to judge that lottery propositions can be justifiably believed on the basis of statistical evidence.

> Just as "belief that *p*" is loosely used to denote both belief that *p strictly speaking* and mere belief that probably *p*, "justified belief that *p*" is loosely used to denote both justified belief that *p* strictly speaking and mere justified belief that probably *p*. This is a usage beyond reproach on many occasions for philosophers and non-philosophers alike … If one draws no distinction between categorical beliefs and their probabilistic counterparts, *of course* I am justified in believing that my lottery ticket will not win. (Sutton, 2007, p. 65)

If these suggestions are on the right track, then one would expect subjects to exhibit a reduced tendency to judge that lottery propositions can be justifiably believed on the basis of statistical evidence when the distinction between probabilistic and categorical lottery propositions is made salient. Our own class room experience (PAE and MS) tentatively offered some anecdotal "evidence" in favor of the hypothesis that people will revise their initial judgment once they are aware of the relevant distinction. This thus leads to the following hypothesis:

Hypothesis 5: *Priming*
People will be less willing to ascribe justified belief in lottery propositions when primed with a question about probabilistic lottery propositions. This effect will be particularly pronounced when the evidence is statistical.

Two further hypotheses are concerned with the following claim that is often made in combination with the lottery intuition – and here again the intention would seem to be that the claim has appeal for philosophers and lay people alike:

> *No matter how high the odds* that the ticket will not win, it strikes us that the ticket-holder doesn't *know* that his ticket will not win. (Vogel, 1990, p. 16, emphasis in the original, bold added)

Similarly, Cohen writes:

> Suppose S holds a ticket in a fair lottery of n tickets, where the probability 1 - 1/n of losing is very high. Does S know that his ticket will lose? Although (if *n* is suitably large) S has good reasons to believe that he will lose, it does not seem right to say that S knows that he will lose. This remains true *for arbitrarily large n.* (Cohen, 1988, p. 92, emphasis added)

Finally, according to Williamson:

> On the merely probabilistic grounds that your ticket was only one of very many, I assert to you flat-out "Your ticket did not win", without telling you my grounds. Intuitively, my grounds are quite inadequate for that outright unqualified assertion, *even though one can construct the example to make its probability on my evidence as high as one likes, short of 1, by increasing the number of tickets in the lottery.* You will still be entitled to feel some resentment when you later discover the merely probabilistic grounds for my assertions. (Williamson, 2000, p. 246, emphasis added)

These authors, and others, share the judgment that one cannot know a lottery proposition on the basis of purely probabilistic evidence *no matter how strong that evidence is* and no matter how probable the lottery proposition is made.

However, in a follow-up study to Wells' influential paper mentioned above, Wright, Maceachern, Stoffer, & Macdonald, 1996) showed that the Wells effect does become much less pronounced as the relevant probabilities are increased. Wright and colleagues showed that willingness among mock jurors to make liability judgments on the basis of statistical evidence increased significantly when the probabilistic strength of that evidence was substantially increased. Hence, Wright and colleagues, as well as our own classroom experience, which tentatively suggested that students confronted with extreme odds tended to change their initial judgment about whether they can know a lottery proposition, provides reason to test whether knowledge judgments are, after all, affected by an increase in the relevant probability. Hence, we tested the following hypothesis:

> ***Hypothesis 6:*** *Large Lotteries*
> When the evidence is statistical, people will have a stronger tendency to ascribe knowledge of lottery propositions if the lottery in question is very large.

Lastly, given the vast number of studies in psychology on so-called anchoring effects, we decided to also test whether the lottery intuition was subject to such effects.

> ***Hypothesis 7:*** *Anchoring*
> When the evidence is statistical, people's willingness to ascribe knowledge of lottery propositions will be sensitive to the size of the lotteries they have been asked about previously. More precisely, when presented with a sequence of increasing or decreasing lotteries, people's willingness to ascribe knowledge of lottery propositions will be subject to an anchoring effect.

It's important to point out that this last hypothesis presupposes, in effect, that the Large-Lotteries hypothesis is borne out, at least to some extent. That is, the lottery intuition won't be a serious candidate for an anchoring effect unless it does exhibit some sensitivity to the size of the lottery involved.

3. Methods

3.1. General survey design and recruitment

We recruited participants for our surveys using Amazon Mechanical Turk. Participants were paid between \$0.25–\$0.50, depending on survey (they were not allowed to take part in more than one survey). We restricted participation to native English speakers only. Participants were asked for personal information (age, native language, philosophy background, and highest completed degree). Following a suggestion from Aust, Diedenhofen, Ullrich, and Musch (2013), participants were offered the opportunity, after having filled out the survey, to note whether they'd done so seriously without affecting their pay. We added comprehension questions and deleted the three fastest and three slowest responses.[8] In total, we recruited 500 participants for 2 different surveys: 300 for survey 1; 200 for survey 2. Each survey had two versions and we recruited 200 for survey 1 version 1 and 100 for the others. Numbers of valid participants for survey 1: 182 for version 1 and 87 for version 2. Survey 2: 80 for version 1 and 88 for version 2. The reason for doubling the numbers for version 1 was that subjects selected their responses from drop-down menus in our surveys and we wanted to test for ordering effects, that is, whether the order in which responses appear in the dropdown box made a difference. No ordering effect was detected, and so the results were pooled.

In our vignettes we made the relevant probabilities of winning/losing explicit and we presented them in terms of ratios as well as percentages. Also, we emphasized that a winning ticket must match *all* lottery numbers, and that *all* other tickets are considered losing tickets.[9] In order to contrast *belief* explicitly with *justified belief* or *knowledge*, we built into the vignettes the assumption that the subject believes that the ticket has lost.[10] More generally, we first asked for Yes/No answers to the justification/knowledge question, followed by a question about confidence in the relevant answer using a five-point Likert scale, with the options: totally confident, highly confident, moderately confident, somewhat confident, not at all confident.

Lastly, the data-sets and code supporting our results are available in the Zenodo repository in a citable format (https://doi.org/10.5281/zenodo.815622) and on GitHub (https://github.com/iandurbach/lotteryjudgements). Any updates will be posted on GitHub.

3.1.1. Survey 1

In order to test for hypothesis 2, we introduced the following simple vignette:

Vignette 1: Probabilistic evidence

Harry owns a single ticket in a lottery. His numbers are 23-42-12-8-28-31. The only way for a ticket to win is to match all six numbers, and any other ticket is a losing ticket. Harry wasn't able to watch the evening news at which the winning numbers were announced. He recalls from his statistics class that the chance of a single ticket winning this lottery is one in 14 million, i.e. there is a 99.999993% chance that his ticket has lost.

Suppose Harry believes that the ticket *has* lost.
Is Harry justified in believing that the ticket has lost? (Yes/No – how confident are you about your answer?)

In order to test for hypothesis 3, we offered the following further 2 vignettes:

Vignette 2: Testimonial evidence
Harry owns a single ticket in a lottery. His numbers are 23-42-12-8-28-31. The only way for a ticket to win is to match all six numbers, and any other ticket is a losing ticket. Harry has just watched the evening news at which the winning numbers were announced and none of his numbers match.

Suppose Harry believes that the ticket *has* lost.
Is Harry justified in believing that the ticket has lost?

Taking again our cue from Turri and Friedman, we also introduced so-called mixed cases in which testimonial evidence is presented along with statistical evidence pertaining to its reliability.

Vignette 3: Mixed evidence
Harry owns a single ticket in a lottery. His numbers are 23-42-12-8-28-31. The only way for a ticket to win is to match all six numbers, and any other ticket is a losing ticket. Harry has just watched the evening news at which the winning numbers were announced and none of his numbers match. The following TV show, however, features TV blunders and reports on a TV anchor who once read out the wrong results (not a single number the anchor reported was correct). A statistician reliably calculates the chance of this happening is one in a million, i.e. there is a 99.9999% chance that the anchor read out the correct results.

Suppose Harry believes that the ticket *has* lost.
Is Harry justified in believing that the ticket has lost?

In order to test for hypotheses 4 and 5, we created a slightly revised version of survey 1, which makes the distinction between a standard and probabilistic lottery proposition salient but is otherwise the same. To wit, the first vignette was adjusted as follows.

Vignette 1: Probabilistic evidence*
Harry and Jennifer together own a single ticket in a lottery. Their numbers are 23-42-12-8-28-31. The only way for a ticket to win is to match all six numbers, and any other ticket is a losing ticket. Neither Harry nor Jennifer were able to watch the evening news at which the winning numbers were announced. They recall from their statistics class that the chance of a single ticket winning this lottery is one in 14 million, i.e. there is a 99.999993% chance that the ticket has lost.

Suppose Harry believes that the ticket is *very likely* to have lost.
Is Harry justified in believing that the ticket is very likely to have lost?

Suppose that in the previous question (Q1) Jennifer believes that the ticket *has* lost.
Is Jennifer justified in believing that the ticket has lost?

3.1.2. Survey 2

The second survey was intended to test hypotheses 1, 6, and 7 (while also providing further data for hypotheses 2 and 3). In contrast to survey 1, we specified that the ticket is indeed a losing ticket, we specified the prize-money involved and kept that winning amount – the stakes – fixed throughout our vignettes. We then tested for three phenomena: first, whether subjects are willing to ascribe knowledge of a lottery proposition, second, whether the strength of the statistical evidence has effects on a subject's willingness to ascribe knowledge (or justification), and lastly whether subjects offer different answers with regards to the knowledge/justification question if they have been previously exposed to larger or smaller lotteries.

We used the following vignettes:

> *Vignette 1 (survey 2)*
> Harry owns a single ticket in a local raffle. His number is 9464. There is at most one winning ticket. The winner receives 100,000 US Dollars. Harry wasn't able to watch the local news at which the winning ticket was announced but he knows that the chance of winning is 1 in 10,000, i.e. there is a 99.99% chance that his ticket has lost.
>
> Suppose Harry believes the ticket has lost and suppose that his ticket is a losing ticket. Is Harry justified in believing that the ticket has lost? (Yes/No – how confident are you about your answer?)
>
> With regard to the previous lottery in Q1, does Harry know that the ticket has lost? (Yes/No – how confident are you about your answer?)

We asked subjects the same question but changed the relevant probabilities involved to 1 in 14 Million, and 1 in 100 Million.

A second group of subjects were exposed to similar vignettes, however, involving first the *Universal Lottery* (with a chance of winning of 1 in 10^{80}), followed by a reduction of probabilities, to 1 in 100 Million chance of winning, and 1 in 14 Million.

Lastly, similarly to survey 1, we also tested for so-called mixed cases with respect to the knowledge and justification question. To do so, we introduced, as the last vignette in survey 2, the so-called Mafia case (very similar to a case used in Turri & Friedman, 2014; Friedman & Turri, 2015), which was then compared to the relevant vignette using the same probability of winning.

> *Vignette Mafia*
> Harry owns a single ticket in this week's lottery. His numbers are 23-42-12-8-28-31. The only winning ticket is one that matches all six numbers, and all others are losing tickets. The winner receives 100,000 US Dollars. He recalls from his statistics class that the chance of a single ticket winning this lottery is one in 14 million, i.e. there is a 99.999993% chance that his ticket has lost.
>
> Harry hasn't seen the official results of the lottery. However, he has just read a reliable news report that the Mafia has in fact rigged the latest lottery and thus the chance of any one ticket winning the lottery is now 1 in 100 Million, i.e. there is now a 99.999999% chance that his ticket has lost.

3.2. Statistical methods: General remarks

The statistical methodology that we employ is not widely used in philosophy. We provide a general motivation for it here, before explaining how it was implemented in analyzing the data from each survey. Readers impatient for the results can move directly to Section 4 without loss of continuity. Our methodology contrasts with a more conventional approach in which each hypothesis is assessed in isolation using an independent statistical model (in the form of a formal hypothesis test). The conventional approach has the advantage that relatively simple statistical models can be used. As each hypothesis is treated "on its own," there is typically a single outcome of interest and a single explanatory variable. For example, our hypothesis that people will be less willing to ascribe justified belief in a lottery proposition when primed with a probabilistic question would be assessed by forming two groups of responses, one containing all primed responses and another containing all those obtained without priming. We would then calculate the proportion ascribing justified belief in each group, and compare these two proportions formally using a t-test.

This approach, however, suffers from three drawbacks. Firstly, as the number of independent tests increases, the probability of finding a statistically significant result in error – a false positive or type-I error – increases. Various adjustments can be made to the level at which "significance" is flagged but these are not without problems of their own.[11] Secondly, because the same participant answers multiple questions, a group can contain more than one response from the same participant, violating the assumption that observations are independent. For example, consider again our priming hypothesis: the same participant is primed (or not), and then answers several questions ascribing justified belief based on different types of evidence. Now, if we include all his or her answers in the primed group of responses – pooling across evidence type – we violate the assumption that responses are independent. If we instead choose to separate out different evidence types, evaluating the priming hypothesis independently in each, the number of hypotheses proliferates, and with it the associated probability of a type-I error. Finally, testing a hypothesis in isolation by definition focuses on the effect of one explanatory variable while ignoring others, which can lead to spurious results when explanatory variables combine with one another to influence the outcome, as is often the case.

Modern statistical practice thus strongly advocates modeling the effect of all relevant explanatory variables on an outcome jointly in a single model, and accounting for correlations introduced by multiple responses from the same participant using an appropriate method. By far the two most dominant methodologies in this area are generalized linear mixed models (GLMMs; see, e.g., McCulloch & Neuhaus, 2001) and generalized estimating equations (GEEs; see, e.g., Hardin, 2005). Both of these model the joint effect of a set of explanatory variables on an outcome of interest, thus addressing the first and third of the drawbacks mentioned above, but differ in the way in which they model correlations introduced by dependent observations.

A detailed assessment of these differences is beyond the scope of our paper.[12] But, to summarize, a feature of dependent observations is that each participant has his or her own "effect" (by virtue of answering multiple questions). GLMMs model these explicitly, by assuming that participant effects are drawn from an appropriately chosen distribution. This results in the estimation of an effect for each participant in the study, which is sometimes useful, for example, where one wishes to make predictions for individual participants. The effects of other variables in the model, say of priming, are interpreted as conditional on a participant, that is, the effect of priming if one keeps the participant fixed. GEEs make weaker assumptions about the nature of the participant effects. They control for, but do not explicitly estimate, the effect of each participant. The effects of other variables in the model are interpreted as averaged over all participants in the study. GEEs are generally considered more appropriate when interest is on average effects across the entire sample of participants, rather than on the effect that individual participants have, in which case GLMMs are preferred.[13] Our interest is almost exclusively on the effect of our main experimental variables (different type of evidence, priming, anchoring), and while the dependencies induced by the presence of multiple responses per participant should be accounted for, they are not, for us, of primary interest. We thus prefer GEEs, which we implement using the *geepack* package in R.[14]

3.2.1. Statistical methods: Survey 1

In survey 1, the outcome – the number of positive responses to each justification question – was assumed to be binomially distributed, conditional on the probability of holding the relevant belief to be justified. Potential correlations between answers for the same respondent were accounted for using generalized estimating equations with respondents treated as the subject variable and independent correlation matrices used. In our first analysis, the probability of holding belief in the lottery proposition to be justified was modeled as being dependent on the type of evidence used (statistical, testimonial, or mixed) and on whether respondents were first primed with the question about the probabilistic lottery proposition. The explanatory variables thus include categorical fixed effects for type of evidence (three levels) and for priming (two levels), and interaction terms between these two effects. As usual for binomial GEEs, the linear dependency between the probability of a justified belief attribution and the explanatory variables was specified using a logit link function. Significance was assessed using Wald chi-square statistics. Following model fitting, *post hoc* tests were used to compare means involving primed and unprimed responses within each evidence type (i.e., statistical, testimonial, or mixed) and to compare means involving different evidence types within each of the primed and unprimed subsets of responses. Means involving different types of evidence and different priming conditions were not compared. The Sidak method was used to adjust for multiple comparisons in these *post hoc* tests.

In our second analysis, the probability of a justified belief ascription was modeled as being dependent on the type of evidence provided (statistical, testimonial, or mixed) and on whether the question was about the lottery proposition or the probabilistic lottery proposition. The explanatory variables thus included categorical fixed effects for type of evidence (three levels) and for the proposition involved (two levels), and interaction terms between these two effects. In other respects, the analysis is conducted as described above.

Confidence responses were strongly skewed toward the high end of the confidence scale. As a result, we transformed confidence responses into a binary scale (totally confident/not totally confident) and analyzed the resulting transformed variable using binomial GEEs, as described above.

3.2.2. Statistical methods: Survey 2

Data were again analyzed using generalized estimating equations (GEEs) implemented in R. The number of positive responses to each justification question was assumed to be binomially distributed, conditional on the probability of judging the relevant belief to be justified. Potential correlations between answers for the same respondent were accounted for using generalized estimating equations with respondents treated as the subject variable and independent correlation matrices used. In our first analysis, the probability of judging the belief in the lottery proposition to be justified was modeled as being dependent on the type of evidence (statistical, testimonial, or mixed) and on whether respondents were first primed with the *small lottery* (1 in 10,000 chance of winning) or with the *universal lottery* (1 in 10^{80} chance of winning). The explanatory variables thus include categorical fixed effects for type of evidence (three levels) and for priming (two levels), and interaction terms between these two effects. In our second analysis, the probability of a positive response to the effect that the lottery proposition is known was modeled as being dependent on the type of evidence (statistical, testimonial, or mixed) and on whether the small lottery or universal lottery anchor was used. The explanatory variables thus include categorical fixed effects for type of evidence (three levels) and for priming (two levels), and interaction terms between these two effects. In other respects, the analysis was conducted as described for survey 1. Confidence responses were again transformed into a binary scale (totally confident/not totally confidence) and analyzed using binomial GEEs, as for survey 1.

4. Results

In this section, we present a detailed analysis of the survey results. Figures are used as descriptive summaries of responses, while significant results are summarized for each survey in the respective table. In Section 5, we then discuss how these results affect the seven hypotheses presented in Section 2.

4.1. Survey 1: Main results

We can summarize four main results from survey 1 as follows (see Table 1 and Figure 1):

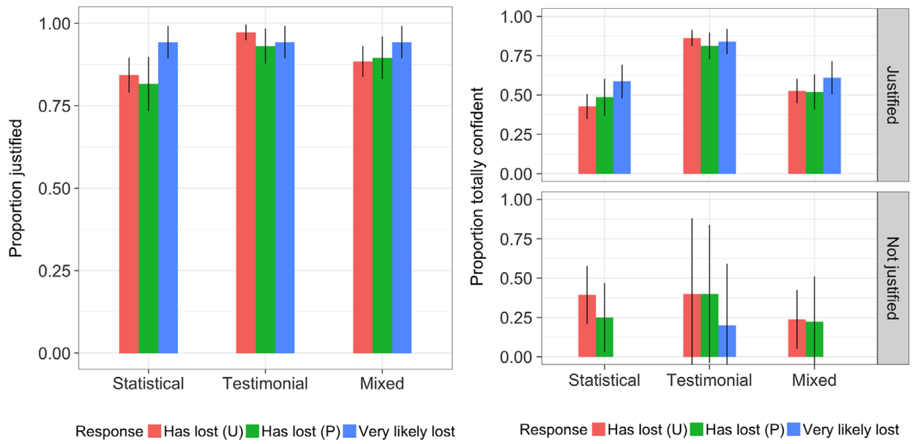

Figure 1. Tendency to ascribe justified belief in the lottery proposition depends on type of evidence used, but not on priming (U = unprimed, P = primed).
Notes: The left-hand plot shows the proportion supporting justified belief in the lottery proposition while the right-hand plot shows the proportion reporting total confidence in their judgments. Bars indicate 95% confidence intervals (wide confidence intervals are due to low number of responses in that group).

Table 1. Effect sizes and summaries obtained from statistical models fitted to the data from survey 1.

Outcome	Predictor	Predictor levels	Wald χ^2	DoF	*p*-value	Significantly different pairs
Response (J/not-J)	Type of evidence	(S)tatistical, (T)estimonial, (M)ixed	28.9	2	<0.001	S & T, T & M
	Priming	(U)nprimed, (P)rimed	0.42	1	0.520	
Response (J/not-J)	Type of evidence	Statistical, testimonial, mixed	22.4	2	<0.001	S & T
	Proposition	LP (U/P), PLP	4.16	1	0.041	LP & PLP
	Proposition × Evidence		6.11	2	0.047	See text
Confidence	Justified belief	Justified, not justified	21.2	1	<0.001	J & not-J
	Type of evidence	Statistical, testimonial, mixed	91.5	2	<0.001	S & T, T & M
	Priming	(U)nprimed, (P)rimed	0.1	1	0.802	
	Priming × Evidence		6.4	2	0.041	See text
Confidence	Justified belief	Justified, not justified	27.7	1	<0.001	J & not-J
	Type of evidence	Statistical, Testimonial, Mixed	74	2	<0.001	S & T, T & M
	Proposition	LP (U/P), PLP	7.5	2	0.024	LP & PLP
	Just × Evidence		7.6	2	0.022	See text
	Just × Proposition		5.6	1	0.018	See text

1. The vast majority of participants agree that belief in lottery propositions is *justified*. Participants who agreed that belief in lottery propositions is justified were also more confident in their answer than those opposing it.
2. Participants were more inclined to ascribe justified belief in a lottery proposition on the basis of testimonial evidence, in contrast to the statistical and mixed evidence types. They also exhibited greater confidence in their answers.
3. Participants were more inclined to ascribe justified belief in a probabilistic lottery proposition than a categorical lottery proposition. Greater confidence was also reported when ascribing justified belief of probabilistic lottery propositions. Participants' tendency to ascribe justified belief in a probabilistic lottery proposition was consistent across the testimonial, statistical, and mixed evidence types.
4. Priming, by first presenting a probabilistic lottery proposition, exerted no significant effect over subsequent judgments about whether belief in a lottery proposition was justified, or over confidence in that judgment.

4.1.1. Result 1: The vast majority of participants agree that belief in lottery propositions is justified

Pooling all responses, and thus disregarding the types of evidence on the basis of which subjects made their judgments, around 89% of respondents answered that Harry/Jennifer were both justified in believing that the ticket had lost (715 of 800 responses).

Subjects who supported justified belief in LP tended to be far more confident in their response than those who did not. Of 711 responses in support of justified belief in the lottery proposition (4 subjects did not report confidence ratings), 437 of these (62%) indicated *total* confidence in this response. Among the 84 valid responses against justified belief in lottery propositions, only 26 (31%) were totally confident. The effect of judging the lottery proposition to be justifiably believed on reported confidence level is highly significant ($z = 4.18$, $p < 0.001$).

4.1.2. Result 2: Participants were more inclined to ascribe justified belief in a lottery proposition on the basis of testimonial evidence, in contrast to the statistical and mixed evidence types

Subjects judging on the basis of testimonial evidence were relatively more likely to judge that the lottery proposition is justifiably believed (96%, compared to 83% and 89% of those seeing statistical and mixed evidence, respectively, $\chi^2 = 91.5$, $p < 0.01$). Post hoc tests indicated that the testimonial group significantly differed from both statistical and mixed groups ($z = 5.00$, $p < 0.001$ and $z = 2.90$, $p = 0.011$, respectively), but that the statistical and mixed groups were only marginally different ($z = 2.32$, $p = 0.060$).

Confidence was significantly higher when the evidence was testimonial, but only among those supporting justified belief in the lottery proposition ($\chi^2 = 6.4$,

$p = 0.041$). 85% of respondents supporting justified belief reported total confidence when receiving testimonial information, compared to 44 and 52% when using statistical and mixed information, respectively. Among those not supporting justified belief in the lottery proposition, 40% of respondents reported total confidence when receiving testimonial information, compared to 34 and 23% when using statistical and mixed information, respectively. Post hoc tests indicated that, among supporting responses, those obtained by testimonial information differed significantly from both statistical and mixed groups ($z = 9.51$, $p < 0.001$ and $z = 8.47$, $p < 0.001$, respectively), but that none of the differences among responses opposing justified belief were significant (all $z < 1.40$, all $p > 0.163$).

4.1.3. Result 3: Participants were more inclined to ascribe justified belief in a probabilistic lottery proposition than a categorical lottery proposition

Although subjects were asked about Harry's belief in a probabilistic lottery proposition primarily to test for the presence of a priming effect, it is instructive to note that subjects answered this question differently than they answered questions about categorical lottery propositions, which suggests that subjects understood the difference between PLP and LP. More subjects supported justified belief in the probabilistic lottery proposition (246 of 261 responses, 94%) than the categorical lottery proposition (Result 1, 89%). This difference is moderately significant ($\chi^2 = 4.16$, $p = 0.041$). In contrast to the case of categorical lottery propositions, participants' willingness to ascribe justification for the probabilistic lottery proposition was independent of the type of evidence presented: precisely the same proportion of subjects, 94%, judged that one could justifiably believe the probabilistic lottery proposition using any of the statistical, testimonial, or mixed evidence formats. This interaction effect, between type of evidence and the type of proposition being probed, is also moderately significant $\chi^2 = 6.11$, $p = 0.047$).

As reported in Result 1, subjects who supported justified belief in the lottery proposition tended to be more confident in their response than those who did not. This effect is stronger for subjects supporting justified belief in the probabilistic lottery proposition ($\chi^2 = 5.56$, $p < 0.001$). 68% of those who supported justified belief in the probabilistic lottery proposition indicated total confidence, compared to just 7% of those who opposed it. Recall that these proportions were 62 and 33% for the standard lottery proposition – smaller but still significant ($z = 3.98$, $p < 0.001$).

4.1.4. Result 4: We found no significant effects indicating priming. First presenting a probabilistic lottery proposition exerted no statistically recoverable effect over subsequent judgments about whether belief in a lottery proposition was justified, or over confidence in that judgment

Similar proportions of subjects supported justified belief in the lottery proposition in the primed and unprimed groups. Of 260 primed responses, 229 (88%)

supported justified belief, compared with 486 of 540 unprimed responses (90%). This difference is not significant ($z = 1.02$, $p = 0.31$). Confidence ratings were also similar across the two groups: 61% of primed subjects supporting justified belief in the lottery proposition reported total confidence, compared with 62% for un-primed subjects. Confidence was lower, but still similar across priming conditions when justified belief was opposed: 27 and 33% of primed and unprimed responses were totally confident, respectively. None of these differences proved statistically significant ($z = 0.03$, $p = 0.98$; $z = 0.71$, $p = 0.48$, respectively).

The effect of the type of evidence was also robust with respect to priming, being clearly demonstrated in both primed and unprimed groups, but with no significant difference in the size of the effect between groups. The proportion of primed and unprimed subjects supporting justified belief in the lottery proposition was, respectively, 93 and 97% for testimonial information, 82 and 84% for statistical information, and 90 and 88% for mixed information. No significant interaction between the type of evidence involved and priming treatment was found ($\chi^2 = 2.47$, $p = 0.29$). Similarly, greater confidence was reported when using testimonial information, regardless of priming. Among primed subjects, total confidence was reported by 79% in the testimonial group, and 44 and 50% in the statistical and mixed groups, respectively. Among unprimed subjects, the proportions were 85, 42, and 49%, respectively. Again, no significant difference in the effect of evidence type across priming treatments was found ($\chi^2 = 1.3$, $p = 0.54$).[15]

4.2. Survey 2: Main results

While survey 2 is primarily designed to test for anchoring effects as well as questions regarding people's willingness to ascribe knowledge in these cases, it overlaps with survey 1 in that it can also be used to gauge overall levels of support for justified belief in the lottery proposition, and confidence in those judgments. Again, pooling all the results and disregarding for a moment the type of evidence on which the judgments are made, we find nearly identical results to those reported for survey 1, corroborating our result 1 above: 442 of 493 (90%) of responses in survey 2 answered that Harry was justified in believing that he had lost, compared with 89% in survey 1. Overall confidence was lower than in survey 1, but respondents who supported justified belief in the lottery proposition again tended to be far more confident. Of 439 responses in support of justified belief in the lottery proposition (3 subjects did not report confidence ratings), 183 of these (42%) indicated total confidence, compared to 62% in survey 1. Note that this discrepancy in confidence may simply be because survey 2 did not contain a vignette using testimonial evidence, while survey 1 did, and it is here that the confidence values were found to be very high (compare 4.1.2).

Among the 51 valid responses that denied justified belief in the lottery proposition, 10 (20%) were totally confident, compared to 31% in survey 1. The effect of positive or negative response on reported confidence is again significant ($\chi^2 = 5.58$,

$p = 0.018$) though not to the extent reported in survey 1 ($\chi^2 = 27.7$, $p < 0.001$). Besides reinforcing overall support for people's tendency to ascribe justified belief in lottery propositions, our main results from survey 2 are as follows (see Table 2 and Figure 2 unless otherwise stated):

1. The majority of participants denied that lottery propositions are known. Similar confidence was reported when ascribing either justified belief or knowledge, but less confidence was reported when denying justified belief than when denying knowledge.
2. Participants were more inclined to ascribe both justified belief and knowledge of lottery propositions, and were more confident in their responses, when presented with the additional contextual information of the so-called Mafia case.

Table 2. Effect sizes and summaries obtained from statistical models fitted to the data from survey 2.

Outcome	Predictor	Predictor Levels	Wald χ^2	DoF	p-value	Significantly different pairs
Response (J/not-J); (K/not-K)	Type of evidence	1:14 M, 1:100 M, Mafia	12.3	1	0.002	1:14 M & Mafia, 1:100 M & Mafia
	Epistemic status	JB, knowledge	152.5	1	<0.001	JB & K
	Anchor	Small, Universal	0.2	2	0.623	
	Anchor x Evidence		4.7	1	0.093	See text
	Anchor × Epistemic Status		2.8	1	0.097	See text
Confidence	Response	J/K, not-J/not-K	5.58	1	0.018	J/K & not-J/not-K
	Type of evidence	1:100 M, Mafia	7.33	2	0.26	1:14 M, Mafia
	Epistemic Status	JB, Knowledge	5.62	1	0.018	JB & K (for not-J/not-K only)
	Anchor	1:10 K, Universal	0.07	1	0.789	
	Response × Epistemic Status		6.54	1	0.011	See text

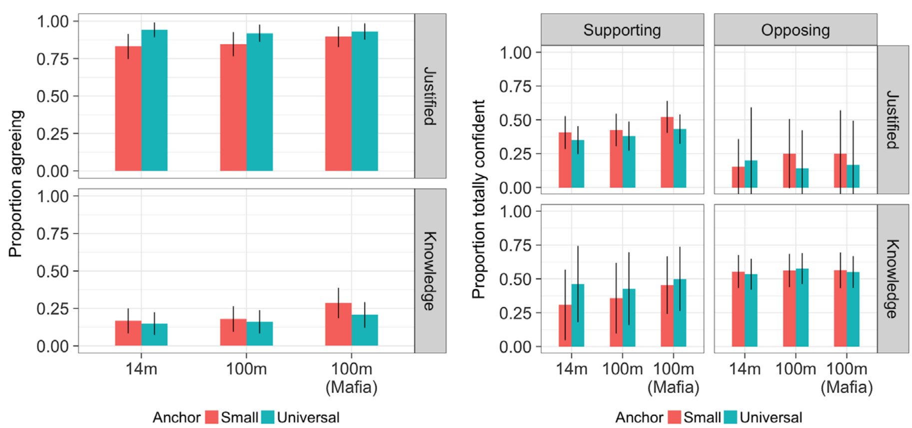

Figure 2. Justified belief in LP is more often supported than knowledge of LP.
Notes: The left-hand plot shows the proportion supporting justified belief in, or knowledge of, LP while the right-hand plot shows the proportion reporting total confidence in their judgments. Bars indicate 95% confidence intervals (wide confidence intervals are due to low number of responses in that group).

Table 3. Differences in judgment with respect to small lottery and universal lottery.

Outcome	Predictor	Predictor levels	Wald χ^2	DoF	*p*-value	Significantly different pairs
Response (J/not-J); (K/not-K)	Size	Small, universal	5.6	1	0.018	Small & Universal (JB only)
	Epistemic status	JB, knowledge	142.2	1	<0.001	JB & K
	Size × Epistemic Status		4.3	1	0.039	See text
Confidence	Response	J/K, not-J/not-K	0.68	1	0.408	
	Epistemic status	JB, knowledge	13.50	1	<0.001	JB & K (Small only)
	Size	Small, universal	3.78	1	0.052	Small & Universal (JB only)
	Size × Epistemic Status		7.36	1	0.007	See text
	Response × Epistemic Status		4.10	1	0.043	See text

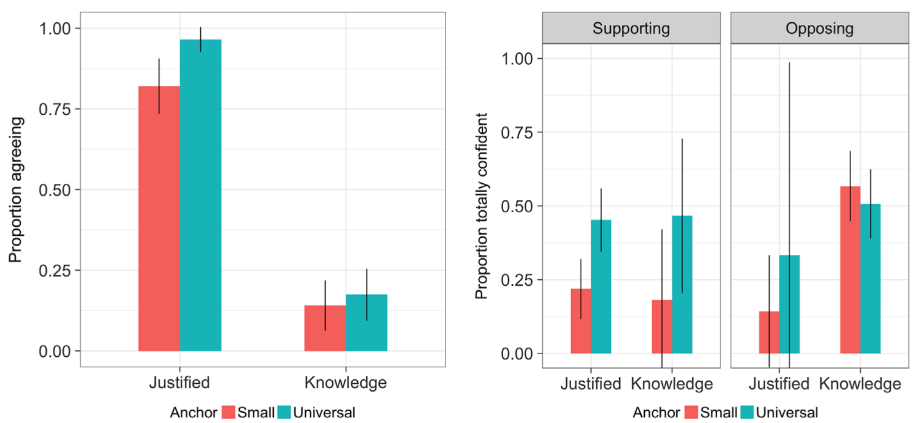

Figure 3. Differences in judgment with respect to small lottery and universal lottery.
Note: Bars indicate 95% confidence intervals (wide confidence intervals are due to low number of responses in that group).

3. Changing the size of the lottery, and the underlying probabilities, had no effect on participants' willingness to ascribe knowledge, but did have an effect on the participants' willingness to ascribe justified belief (Table 3, Figure 3).
4. Anchoring, by first asking participants a question about a larger (i.e., universal lottery) or smaller lottery, exerted no effect over subsequent judgments about justified belief or knowledge of lottery propositions, or over confidence in those judgments.

4.2.1. *Result 1: The majority of participants denied that lottery propositions are known*

Many more respondents supported justified belief in lottery propositions (442 of 493 responses, 90%) than knowledge in lottery propositions (94 of 494 responses, 19%). This difference is highly significant ($\chi^2 = 152.5$, $p < 0.001$).

There is some very weak evidence to suggest that the difference observed between support for justified belief and knowledge depends on the anchoring question used ($\chi^2 = 2.8$, $p = 0.097$). Support for justified belief increased when the Universal lottery was used as anchor (93 vs. 86% using the small lottery anchor). However, support for knowledge of the same marginally *decreased* with the Universal lottery as anchor (17 vs. 21% using the small lottery as anchor).

Post hoc tests showed the differences between knowledge ascription and justified belief ascription within each anchoring group to be highly significant (for the small lottery anchor $z = 8.21$, $p < 0.001$; for the Universal anchor $z = 8.88$, $p < 0.001$), with the increased significance seen for the small lottery anchor providing the basis for the marginally significant interaction. Given its marginal significance level, this interaction effect must be interpreted with caution, but we report it here for completeness and as a possible avenue for future work.

Confidence was significantly higher when concerned with knowledge of lottery propositions than justified belief ($\chi^2 = 5.62$, $p = 0.018$), but only among those *opposing* the relevant statement ($\chi^2 = 6.54$, $p = 0.011$). Among all opposing or disagreeing responses, 20% reported total confidence when the statement was about justified belief, compared to 56% when the statement was about knowledge. No such difference was observed among supporting responses, where 42 and 43% reported total confidence when the statement was about justified belief and knowledge, respectively.

Post hoc tests confirmed that the difference between groups opposing justified belief and knowledge was significant ($z = 2.85$, $p = 0.026$). Thus, not only did many more subjects reject the claim that Harry could know he had lost, but they were also relatively more confident in this rejection.

4.2.2. Result 2: Participants were more inclined to ascribe both justified belief and knowledge of lottery propositions, when presented with additional contextual information of the 'Mafia' case

Subjects were relatively more likely to support justified belief and knowledge statements when given the additional information that the Mafia had manipulated the lottery ($\chi^2 = 7.33$, $p = 0.026$). 92 and 24% of responses supported justified belief and knowledge, respectively, compared to 89 and 17% (1:100 million chance) and 89 and 16% (1:14 million chance) when only statistical information was provided.

Post hoc tests indicated that the Mafia group significantly differed from both 1:100 million (no Mafia) and 1:14 million groups ($z = 3.10$, $p = 0.005$ and $z = 3.29$, $p = 0.003$, respectively), but that the latter two groups are not significantly different ($z = 0.23$, $p = 0.994$). Although empirical proportions show that additional information has a greater impact on support for knowledge (24% in the Mafia group; 17 and 16% in the other groups) than on support for justified belief, this interaction between epistemic status and type of evidence proved not to be significant ($\chi^2 = 3.2$, $p = 0.203$).

A very marginally significant interaction exists between the type of evidence used and the anchoring question used ($\chi^2 = 4.7$, $p = 0.093$). *Post hoc* tests showed that the increased support for both justified belief and knowledge observed using the Mafia case was only significant using the small lottery anchor (comparison with 1:14 Million group, $z = 3.28$, $p = 0.006$; comparison with 1:100 Million group, $z = 2.78$, $p = 0.032$; other $p > 0.53$). Among responses obtained with a small lottery anchor, 90 and 29% were supportive of justified belief and knowledge respectively in the Mafia case, compared to 85 and 18% of those exposed to the 1:100 million (but no Mafia) and 83 and 17% of those exposed to the 1:14 million vignettes. The same proportions among responses obtained with the universal lottery anchor were 93%/21% (Mafia), 92%/16% (1:100 Million), and 94%/15% (1:14 Million); differences between the Mafia and other groups are here somewhat smaller. Given the marginal significance of the interaction effect, these *post hoc* tests and the conclusions drawn from them should be interpreted with caution.

Subjects were more confident in their responses when given additional information about the Mafia. As stated above (result 1), higher confidence was reported for those opposing knowledge than those opposing justified belief, but similar confidence levels were reported by those supporting knowledge/justified belief. Supporters of both justified belief and knowledge did, however, report more confidence in the Mafia case (justified belief: 47% with Mafia information, 40 and 38% using statistical information; knowledge: 48% using Mafia information, 39 and 38% using statistical information). No such difference was observed among opposers (justified belief: 21% with Mafia information, 21 and 17% using statistical information; knowledge: 56% using Mafia information, 57 and 54% using statistical information). Note that only the main effect (of type of evidence on confidence) is moderately significant ($\chi^2 = 7.33$, $p = 0.026$). In contrast to survey 1, the effect of type of evidence on confidence in survey 2 does not depend statistically on any of the other experimental variables, as shown by the absence of any significant interactions involving the type of evidence. Thus, additional inferences about the role of epistemic status and support or opposition for this status must be interpreted as speculative. Finally, *post hoc* tests indicated that only the Mafia and 1:14 million groups differed significantly ($z = 2.82$, $p = 0.014$).

4.2.3. Result 3: We found no significant effect on participants' willingness to ascribe knowledge when changing the size of the lottery, that is, the underlying probabilities. However, changes in the size of the lottery did have an effect on the participants' willingness to ascribe justified belief

A significantly lower proportion of subjects supported justified belief when a small lottery was presented, compared to the universal lottery ($z = 2.57$, $p = 0.023$). In the small lottery, 64/78 (82%) of subjects support justified belief in LP; that proportion increases to 84/87 (97%) in the case of the universal lottery. No effect, however, was observed with regard to the knowledge question ($z = 0.58$, $p = 0.96$). Considering the small and the universal lotteries, respectively, the proportion of subjects supporting knowledge of LP was 11/78 (14%) and 15/86 (17%). Collectively, these proportions describe the significant joint (i.e., interaction) effect of lottery size and epistemic status on participant responses ($\chi^2 = 4.3$, $p = 0.039$).

A significantly greater proportion of subjects reported greater confidence in ascribing justified beliefs in LP in the universal lottery compared to the small lottery ($z = 2.99$, $p = 0.016$). No significant differences in confidence were found between those supporting and those opposing justified belief[16], and we thus pooled results across these groups. 39/87 (45%) of those faced with the universal lottery were confident in their response, compared to 16/78 (21%) of those faced with the small lottery. No such difference in confidence was observed in the case of knowledge ($z = 0.08$, $p > 0.99$), where the relevant proportions were 49 and 51%. The interaction effect between anchor and epistemic status was highly significant ($\chi^2 = 7.36$, $p = 0.007$).

4.2.4. Result 4: Judgments about justified belief or knowledge of lottery propositions, and the confidence with which these judgments were made, did not differ significantly over anchoring conditions

As in survey 1, similar proportions of responses supported justified belief in LP across priming groups. Of 234 responses primed with the small lottery anchor, 199 (86%) and 49 (21%) were supportive of justified belief and knowledge respectively, compared with 243 (93%) and 45 (17%) of 261 responses seeing the universal lottery anchor.

This difference is not significant ($\chi^2 = 0.2$, $p = 0.62$). As mentioned in result 1 above, there is weak evidence to suggest that the effect of type of evidence used, in particular the provision of additional contextual information in the form of the Mafia statement, may not be robust to anchoring effects, with support for justified belief increasing with a universal lottery anchor but support for knowledge of the same decreasing ($\chi^2 = 4.7$, $p = 0.093$). We reiterate that this result should be interpreted as speculative.

Confidence ratings are also similar: 47% of responses primed with the small lottery anchor reported total confidence compared with 46% of those seeing the universal lottery anchor ($\chi^2 = 0.07$, $p = 0.78$). No significant interactions with other experimental variables were reported.

5. Discussion

In this section, we will briefly discuss our results in the context of the seven hypotheses we offered in Section 2.

5.1. Hypotheses 1 and 2: No-knowledge and yes-justification

Both hypotheses received strong support from our survey study. Subjects have a clear tendency to deny knowledge of lottery propositions when the evidence is statistical (81% deny knowledge of LP, 19% ascribe knowledge of LP; compare 4.2.1), while most subjects are willing to ascribe a justified belief in the lottery proposition based purely on statistical evidence (83%[17] ascribe justified belief in LP; compare 4.1.2). In general, subjects exhibit high confidence in their answers. These findings are broadly in line with the studies by Turri and Friedman (2014) and Friedman and Turri (2015).

5.2. Hypothesis 3: Testimonial

Testimonial evidence had a significant effect with respect to an increased willingness to ascribe justified belief compared to pure statistical cases, and higher confidence in the given response (if the answer is positive). As such hypothesis 3 also received strong support from our survey (compare 4.1.2). Our findings are again broadly in line with the previous studies by Turri and Friedman (2014) and Friedman and Turri (2015). The issue about mixed cases will be discussed further in Section 6.

5.3. Hypothesis 4: Probabilistic

Our survey provides support for the hypotheses that subjects have a stronger tendency to ascribe justified belief in probabilistic lottery propositions (94%) than in categorical lottery propositions (89%), irrespective of the type of evidence. Noteworthy here is also that confidence levels are higher in the former than in the latter case. Also, we found that subjects' willingness to ascribe justified belief in PLP is not affected by the relevant kind of evidence (compare 4.1.3). As such, the second part of our hypothesis that the increase is most pronounced when the evidence is statistical is also supported by the results of survey 1, since here we have the highest increase (from 83 to 94%).[18]

5.4. Hypothesis 5: Priming

In contrast to our own expectations, we have not found any significant priming effects. While subjects assess probabilistic and categorical lottery propositions differently, in particular when assessing them in the context of statistical evidence (compare 5.3), being made aware of that distinction had no significant effect on their judgments about categorical lottery propositions (compare 4.1.4). Moreover, no significant differences were found with regard to primed and non-primed subjects' confidence levels. As such, this raises doubts about the correctness of Sutton's suggestions quoted earlier that the general tendency to ascribe justified belief in LP is in part due to a conflation between LP and PLP.

5.5. Hypothesis 6: Large lotteries

In contrast to our own expectations, we found no significant changes in subjects' willingness to ascribe knowledge when strongly varying the strength of statistical evidence. This contrasts with our finding that a subject's willingness to ascribe justified belief in LP is affected by the strength of statistical evidence (compare 4.2.3). Also, the confidence of subjects ascribing justified belief increases when the probability of winning decreases. Hence, this result supports the claim made by Vogel, Cohen, Williamson, and others that the odds of a lottery do not affect knowledge ascription. We will discuss this further below.

5.6. Hypothesis 7: Anchoring

Given that we found no significant variation in subjects' willingness to ascribe knowledge, depending upon the size of the lottery, there were no discernible anchoring effects with respect to knowledge ascription either. We may note that there is some very weak evidence that willingness to ascribe justified belief may be affected by the relevant anchors; however, further research is required to investigate this initial finding.

6. General remarks

Our results confirm the basic findings of Turri and Friedman that lay judgments about when we can know lottery propositions largely accord with those of philosophers. In addition, the main findings from our survey can be summarized as follows:

1. Contextual information plays an important role and may trump probabilistic information when ascribing knowledge.
2. The size of a lottery has no effect on whether people are willing to ascribe knowledge of lottery propositions when the evidence is statistical.
3. We did not detect any priming effects, that is, people primed with a question about probabilistic lottery propositions were just as willing to ascribe justified belief in lottery propositions as those that were not primed by this distinction.

Contextual information, such as that provided in the Mafia case, has a significant effect on people's willingness to ascribe knowledge of a lottery proposition. Here we found a significant increase in knowledge judgments from 16 to 24%, and subjects were more confident in their responses as well. Friedman and Turri (2015) observed a similar pattern, with 35% of subjects ascribing knowledge in a Mafia-type case, compared to only 14% in a pure statistical case. Curiously, Turri and Friedman (2014) found no increased willingness to ascribe knowledge of a lottery proposition in a Mafia-type case, as opposed to a statistical case. Comparison between these studies has to be treated with care, however: variations might have to do with the precise nature of the contextual information and the way in which it is provided.[19] In general, questions about when and how additional contextual information can trigger different knowledge judgments may be a fruitful avenue for further research. While the Mafia information did, in our study, generate a significant effect on knowledge judgments, we found only a very marginally significant difference between the mixed evidence and statistical evidence conditions in the first survey on justified belief. Hence, it may be that contextual information is more relevant to knowledge judgments (and the confidence with which they are made) than to justified belief judgments, though this too will need further investigation.

In addition, we found that lay judgments also conform to the claim made by Williamson, Cohen, and others that the size of a lottery makes no difference to whether lottery propositions can be known on the basis of statistical evidence. It is noteworthy, though, that, while a subject's willingness to judge that a lottery proposition is known seems not to be affected by the size of the lottery and the probability of the proposition, when making judgments about justification, subjects did show some tendency to track or respond to the underlying changes in probabilities. This observation is further strengthened by the fact that subjects' confidence in ascribing justified belief in a lottery proposition increased as the chance of winning decreased – a phenomenon not found in the case of knowledge ascriptions. This result suggests a further difference between people's reasoning about justification and about knowledge.

In hindsight, it would have been valuable to test subjects' judgments about whether lottery propositions can be known or justifiably believed in the case of a very small lottery with a higher chance of winning (say 1 in 5) to get a better sense of when the small proportion (roughly 16%) of subjects who do ascribe knowledge on the basis of statistical evidence, and the large proportion who ascribe justification, tend to switch their judgment. A follow-up study of this kind might also help to shed further light on the relationship between knowledge, justification, and liability judgments.

As we noted previously, in a follow-up study to Wells (1992), Wright and colleagues (1996) showed that a majority of subjects are willing to ascribe liability on the basis of purely statistical evidence, provided that the probability of liability is sufficiently high. The study revealed that while mock jurors (in a civil case) are unwilling to ascribe liability if the bare statistical evidence made liability a mere 80% likely, a majority of mock jurors returned a positive liability judgment when this was increased to 99.9%.[20] Hence, it seems that the relevant probabilities are taken into account by subjects when considering liability (at least in a civil case).[21] This may suggest that liability judgments in the given contexts are closer in one respect to justification judgments than knowledge judgments (when the evidence is purely statistical): subjects are generally sensitive to the strength of statistical evidence in liability and justification judgments, while knowledge judgments appear to exhibit no such sensitivity. Further studies which more directly probe the connections between these three kinds of judgment would be needed to test these speculative suggestions.[22]

Further conclusions about the philosophical significance of these results will be guarded. Insofar as philosophers claim that people are inclined to make certain lottery judgments and use this to support a theory directly, or to establish an adequacy constraint on theories, that claim has to be established by the appropriate means and methods. In that respect, our study supports the claim that people tend to judge in line with the lottery intuition. Of course, there can be doubts whether a survey setup is sufficiently reliable to elicit lay people's intuitions, but we took a number of steps (i.e., adding an understanding question, offering subjects the opportunity to say that they were not serious in their response, as well as deleting the fastest and slowest responses) to try and ensure that our surveys were a suitable tool.[23] Hence, we think the onus is on the critic to say why there should be a difference between lay people's thinking as exhibited in surveys and lay people's intuitions properly understood.

There are, of course, various philosophical theories of knowledge and justification that "clash" with our results – in that judgments we found to be widely endorsed would be false according to the theories in question. This clash, in and of itself, provides no objection to a theory and, in many cases, would come as no surprise to its proponents. As discussed above, there are a number of theories of justification that predict that we lack justification for believing lottery propositions on the basis of statistical evidence. Proponents of these theories, however, generally acknowledge that this is a "counterintuitive" prediction that runs contrary to ordinary judgment and, in response, they tend to offer somewhat speculative hypotheses as to what might be driving the recalcitrant judgments. It is here that our results may have a more direct bearing. As discussed, our findings cast doubt on the hypothesis – put most starkly by Sutton, but hinted at by many others – that people's tendency to judge that one can justifiably believe lottery propositions on the basis of statistical evidence is due to a conflation between lottery propositions and probabilistic lottery propositions. Any alternative suggestion as to the psychological mechanisms behind such judgments will be just as amenable to empirical testing.

Notes

1. Some exceptions include Lycan (2006), Hill and Schechter (2007), and Reed (2010).
2. See in particular Harman (1973), Vogel (1990), DeRose, and Hawthorne (2004).
3. Numerous other philosophers have invoked the lottery intuition in the assertion debate, to name a few: Hinchman (2013) and McKinnon (2013).
4. Turri and Friedman cite Nelkin (2000) and Sutton (2007) as proponents of the view. As noted above, Nelkin and Sutton do claim that one cannot know or justifiably believe lottery propositions on the basis of statistical evidence, but at no point, as far as we are aware, do they suggest that our tendency to make the former judgment is explained by our tendency to make the latter judgment. Both Nelkin and Sutton offer arguments for the claim that one cannot justifiably believe a lottery proposition on the basis of statistical evidence. Here, the claim that one cannot know a lottery proposition on the basis of statistical evidence effectively functions as a premise in

these arguments. Hence, their dialectical strategy is very difficult to square with the justification account. In fact, using the knowledge denial, along with certain bridging principles, to derive the justification denial is a common argumentative trope in this area; see, for example, Bird (2007), Smith (2016, ch. 1), and Ichikawa (2014).

5. Compare here, for example, the so-called *expertise defense*, according to which lay people, in contrast to professional philosophers, lack the relevant skills and expertise that is required when assessing thought experiments in an unbiased way. Their assessment of philosophical thought experiments is thus unreliable and of limited significance. This kind of response was first mooted in Weinberg and colleagues (Weinberg, Nichols, & Stich, 2001) and developed further by numerous other authors, such as Williamson (2004), Ludwig (2007), Kauppinen (2007), Williamson (2009a), Horvath (2010), and Williamson (2011). For recent responses and further discussion, see Weinberg, Gonnerman, Buckner, and Alexander (2010), Clarke (2013), Turri (2013), Rini (2014), Mizrahi (2015), Horvath and Wiegmann (2016), and Liao (2016). For a recent survey, see Nado (2014).
6. See, for example, Cappelen (2013); for a response, see Climenhaga (in press).
7. Compare Climenhaga (in press), who offers further examples and discussion of the use of intuitions in epistemology and philosophy more generally.
8. The specific numbers of deletion are as follows. In survey 1 version 1, there was a total of 18 deletions: 6 for time, 7 for failed comprehension, 3 non-english, and 2 for participating in a previous study. In survey 1 version 2, there was a total of 13 deletions: 6 for time, 6 failed comprehension, and 6 for participating in a previous study; some people violated multiple conditions. In survey 2 version 1, there was a total of 20 deletions: 6 for time, 10 failed comprehension, 2 non-english, 5 having a degree in philosophy, and 1 for participating in a previous study. In survey 2 version 2, there was a total of 12 deletions: 6 for time, 1 failed comprehension, 1 non-english, 2 having philosophy degree, and 4 for participating in a previous study.
9. This avoids an ambiguity which may be present in vignettes used in previous studies. After all, in the UK National Lottery, two matching numbers is enough to win a free lucky dip (with close to 1 million tickets matching two numbers each week). Compare: www.national-lottery.co.uk.
10. Compare here, for example, Sripada and Stanley (2012, section 1.2).
11. See for further discussion Bender and Lange (2001).
12. See, however, Gardiner, Luo, and Roman (2009) and Hubbard et al. (2010) for further discussion.
13. Again, for further discussion, consult Hubbard et al. (2010).
14. See Halekoh, Højsgaard, and Yan (2006) and R Core Team (2016).
15. This, of course, isn't to say that there categorically is no priming effect – that would be to accept the null hypotheses, which we like to refrain from doing.
16. Participants who opposed knowledge were marginally more confident than participants who opposed justified belief ($z = 2.05$, $p = 0.079$), while there was no significant difference in the confidence of those supporting knowledge and those supporting justified belief ($z = 0.24$, $p = 0.96$). This is the reason for the significant interaction between response and epistemic status reported in Table 3. What we mean here, however, is that because the size of the anchor plays no *additional* role, pooling across anchoring groups is justified.
17. Note, however, this increases to 97% in the case of the universal lottery. See, for further discussion, 4.2.3 and 5.5.
18. We should note, however, that subjects' willingness to ascribe justified belief in LP seems to track the underlying probabilities. So, while our hypothesis is correct with

respect to the relevant probabilities used in survey 1, the relevant difference may not be as pronounced when the statistical evidence increases in strength. Compare 4.2.3 and 5.5 for further discussion.

19. Indeed, we found a marginally significant interaction in the Mafia case with the small lottery anchor, which may be worth further investigation. Compare Section 4.2.2 for details.
20. While 30% provide a positive verdict in the 80% case, 63% provide a positive verdict in the 99.9% scenario; compare Wright et al. (1996, p. 685).
21. A phenomenon we attempted to recreate – but clearly failed – with regard to knowledge judgments using even higher probabilities.
22. The connection between liability judgments and knowledge judgments has been the subject of recent attention from philosophers. For two somewhat contrasting views see Enoch and colleagues (Enoch, Spectre, & Fisher, 2012) and Blome-Tillmann (2016).
23. The "intuitions" that we (PE and MS) observed in our classrooms could not be recreated in the more controlled survey setting leading us to think that we did unintentionally influence the responses.

Acknowledgments

We would like to thank the audiences of a staff research seminar and a philosophy society talk at the University of Stirling for their comments. In addition, we would like to thank Nora Hanson, Meredith Regan, John Turri, and two anonymous referees of this journal for valuable feedback. Lastly, we like to thank to University of Edinburgh for covering the cost of publication.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

Ian Durbach's work on this paper was supported in part by the National Research Foundation of South Africa [grant numbers 90782, 105782]. Martin Smith's work on this paper was supported by the Arts and Humanities Research Council [grant number: AH/M009610/1].

ORCID

Philip A. Ebert http://orcid.org/0000-0002-3730-0600

References

Aust, F., Diedenhofen, B., Ullrich, S., & Musch, J. (2013). Seriousness checks are useful to improve data validity in online research. *Behavior Research Methods, 45*, 527–535.

Baumann, P. (2004). Lotteries And Contexts. *Erkenntnis, 61*, 415–428.

Bender, R., & Lange, S. (2001). Adjusting for multiple testing–when and how? *Journal of Clinical Epidemiology, 54*, 343–349.

Bird, A. (2007). Judging judging. *Philosophy and Phenomenological Research, 74*, 81–110.

Blome-Tillmann, M. (2016). 'More likely than not'–knowledge first and the role of bare statistical evidence in courts of law. Unpublished manuscript.

Cappelen, H. (2013). *Philosophy without intuitions*. New York, NY: Oxford University Press.

Clarke, S. (2013). Intuitions as evidence, philosophical expertise and the developmental challenge. *Philosophical Papers, 42*, 175–207.

Climenhaga, N. (in press). Intuitions are used as evidence in philosophy. *Mind*. doi:10.1093/mind/fzw032

Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives, 2*, 91–123.

Cohen, S. (1998). Contextualist solutions to epistemological problems: Scepticism, gettier, and the lottery. *Australasian Journal of Philosophy, 76*, 289–306.

Dretske, F. (1970). Epistemic operators. *The Journal of Philosophy, 67*, 1007–1023.

Enoch, D., Spectre, L., & Fisher, T. (2012). Statistical evidence, sensitivity, and the legal value of knowledge. *Philosophy and Public Affairs, 40*, 197–224.

Friedman, O., & Turri, J. (2015). Is probabilistic evidence a source of knowledge? *Cognitive Science, 39*, 1062–1080.

Gardiner, J., Luo, Z., & Roman, L. (2009). Fixed effects, random effects and GEE: What are the differences? *Statistics in Medicine, 28*, 221–239.

Halekoh, U., Højsgaard, S., & Yan, J. (2006). The r package geepack for generalized estimating equations. *Journal of Statistical Software, 15*(2), 1–11.

Hardin, J. (2005). *Generalized estimating equations (GEE)*. Hoboken, NJ: Wiley Online Library.

Harman, G. (1973). *Thought*. Princeton, NJ: Princeton University Press.

Hawthorne, J. (2004). *Knowledge and lotteries*. New York, NY: Oxford University Press.

Hill, C., & Schechter, J. (2007). Hawthorne's lottery puzzle and the nature of belief. *Philosophical Issues, 17*, 102–122.

Hinchman, E. (2013). Assertion, sincerity, and knowledge. *Nous, 47*, 613–646.

Horvath, J. (2010). How (not) to react to experimental philosophy. *Philosophical Psychology, 23*, 447–480.

Horvath, J., & Wiegmann, A. (2016). Intuitive expertise and intuitions about knowledge. *Philosophical Studies, 173*(10), 1–26.

Hubbard, A., Ahern, J., Fleischer, N., Laan, M., Lippman, S., Jewell, N., … Satariano, W. (2010). To gee or not to gee: Comparing population average and mixed models for estimating the associations between neighborhood risk factors and health. *Epidemiology, 21*, 467–474.

Ichikawa, J. (2014). Justification is potential knowledge. *Canadian Journal of Philosophy, 44*, 184–206.

Kauppinen, A. (2007). The rise and fall of experimental philosophy. *Philosophical Explorations, 10*, 95–118.

Liao, S. (2016). Are philosophers good intuition predictors? *Philosophical Psychology, 29*, 1004–1014.

Ludwig, K. (2007). The epistemology of thought experiments: First person versus third person approaches. In P. A. French & H. K. Wettstein (Eds.), *Midwest Studies in Philosophy* (pp. 128–159). New York, NY: Blackwell.

Lycan, W. (2006). On the gettier problem problem. In S. Hetherington (Ed.), *Epistemology Futures* (pp. 148–168). New York, NY: Oxford University Press.

McCulloch, C., & Neuhaus, J. (2001). *Generalized linear mixed models*. Hoboken, NJ: Wiley Online Library.

McKinnon, R. (2013). Lotteries, knowledge, and irrelevant alternatives. *Dialogue, 52*, 523–549.

Mizrahi, M. (2015). Three arguments against the expertise defense. *Metaphilosophy, 46*, 52–64.

Nado, J. (2014). Philosophical Expertise. *Philosophy Compass, 9*, 631–641.

Nelkin, D. (2000). The lottery paradox, knowledge, and rationality. *Philosophical Review, 109*, 373–408.

Nozick, R. (1981). *Philosophical explanations*. Cambridge, MA: Harvard University Press.
Pritchard, D. (2005). *Epistemic luck*. New York, NY: Oxford University Press.
R Core Team. (2016). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing, 2016. Retrieved from https://www.R-project.org/
Reed, B. (2010). A defence of stable invariantism. *Noûs, 44*, 224–244.
Rini, R. (2014). Analogies, moral intuitions, and the expertise defence. *Review of Philosophy and Psychology, 5*, 169–181.
Smith, M. (2010). What else justification could be. *Noûs, 44*, 10–31.
Smith, M. (2016). *Between probability and certainty: What justifies belief*. Oxford: Oxford University Press.
Smithies, D. (2012). The normative role of knowledge. *Noûs, 46*, 265–288.
Sripada, C., & Stanley, J. (2012). Empirical tests of interest-relative invariantism. *Episteme, 9*, 3–26.
Sutton, J. (2007). *Without justification*. Cambridge, MA: MIT Press.
Turri, J. (2013). A conspicuous art: Putting Gettier to the test. *Philosophers' Imprint, 13*(10), 1–16.
Turri, J., & Friedman, O. (2014). Winners and losers in the folk epistemology of lotteries. In J. Beebe (Ed.), *Advances in experimental epistemology* (pp. 45–70). London: Bloomsbury.
Vogel, J. (1990). Are there counterexamples to the closure principle? In G. Ross & M. Roth (Eds.), *Doubting: Contemporary perspectives on scepticism* (pp. 13–27). Dordrecht: Kluwer Academic Publishers.
Weinberg, J., Nichols, S., & Stich, S. (2001). Normativity and epistemic intuitions. *Philosophical Topics, 29*, 429–460.
Weinberg, J., Gonnerman, C., Buckner, C., & Alexander, J. (2010). Are philosophers expert intuiters? *Philosophical Psychology, 23*, 331–355.
Wells, G. (1992). Naked statistical evidence of liability: Is subjective probability enough? *Journal of Personality and Social Psychology, 62*, 739–752.
Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
Williamson, T. (2004). Philosophical 'intuitions' and scepticism about judgement. *Dialectica, 58*, 109–153.
Williamson, T. (2009a). *The philosophy of philosophy*. New York, NY: Blackwell Pub.
Williamson, T. (2009b). Probability and danger. *The Amherst Lecture in Philosophy* (p. 4), Lecture
Williamson, T. (2011). Philosophical expertise and the burden of proof. *Metaphilosophy, 42*, 215–229.
Wright, E., Maceachern, L., Stoffer, E., & Macdonald, N. (1996). Factors affecting the use of naked statistical evidence of liability. *The Journal of Social Psychology, 136*, 677–688.