

Winners and losers in the folk epistemology of lotteries*

JOHN TURRI

john.turri@gmail.com

ORI FRIEDMAN

friedman@uwaterloo.ca

Abstract: Two assumptions anchor most contemporary discussions of knowledge in cases of (large, fair, single-winner) lotteries. First, based on the long odds alone, you don't know that your ticket lost. Second, based on watching a news report of the winning numbers, you do know that your ticket lost. Moreover, it is often treated as an uncontroversial datum that this is how most people view matters. Explaining why people hold this combination of attitudes is then treated as a *criterion* for an acceptable theory of knowledge and knowledge attributions. But do people actually hold the views they're assumed to hold? We did the necessary empirical work to find out. We studied people's reactions to lottery cases and discovered that they respond as predicted. We report those results here. We also evaluate three previous explanations for why people deny knowledge in lottery cases; none of them seems to work. Finally, we present evidence for a new explanation for why some people deny knowledge in lottery cases. We suggest that they deny knowledge in lottery cases due to formulaic expression.

* This is a draft (2013-02-04) of a paper to appear in *Advances in Experimental Epistemology*, ed. James Beebe (Continuum). Comments welcome. Please don't cite, quote or refute without permission. Authorship is co-equal and listed un-alphabetically.

1. Introduction

Suppose that Smith is considering the fate of a particular ticket in a large, fair lottery.¹ After considering the long odds, Smith concludes that the ticket is a loser and, unsurprisingly, Smith is right. Does Smith know that the ticket is a loser, or does he only believe it? As Jonathan Vogel puts it, “No matter how high the odds that the ticket will not win, it strikes us that [Smith] doesn’t *know* that [the] ticket will not win” (Vogel 1990: 292). Call this *the skeptical lottery judgment*.

Now suppose that Brown is considering the fate of that very same lottery ticket. After hearing the winning numbers announced on the nightly news, Brown concludes that the ticket is a loser. Does Brown know that the ticket is a loser, or does she only believe it? As Keith DeRose notes, “after she’s heard the winning numbers announced,” people “judge that [Brown] does know” (DeRose 1996: 570ff). Call this a *nonskeptical lottery judgment*.

This combination of skeptical and nonskeptical tendencies is puzzling. After all, *mistaken testimony* seems much more likely than *winning the lottery*. Indeed, even if you watch the drawing in person and see the winning number with your own eyes, it’s far from clear that *misperception* is any less likely than a *false inference based on the long odds*. Nevertheless, people readily judge that you know the ticket lost after being told the results or watching

¹ Contemporary epistemological discussion of lotteries is vast and traces to Kyburg 1961. Influential recent discussions include Nelkin 2000 and Hawthorne 2004.

the drawing, whereas they readily judge that you don't know after simply calculating the odds.

This paper asks two main questions. *First*, do people actually display the pattern of skeptical and nonskeptical judgment described above? That is, do people share (a) the skeptical lottery judgment in lottery cases involving statistical reasoning and (b) the nonskeptical judgment in lottery cases involving testimony? We find that, yes, people do display this pattern of judgment. *Second*, why do people judge this way? Many explanations have been proposed. We assess three previous proposals and identify a new factor that we believe contributes to skeptical lottery judgment. The new factor is formulaic expression.

It is widely assumed that people conform to the pattern of skeptical and nonskeptical lottery judgment described above. Indeed, theorists claim that it is “uncontroversial” — a “datum” to be explained — that people conform to the pattern (Hawthorne 2004: 8). In short, philosophers assume that they have identified uncontroversial elements of *the folk epistemology of lotteries*. But this sweeping empirical generalization has never been tested. This should inspire caution. For experience shows that armchair predictions often misidentify which epistemological judgments are widespread and uncontroversial. For example, until recently it was widely assumed that virtually everyone shares the intuition that Gettier subjects lack knowledge (Gettier 1963; see Turri 2012a for an overview of the literature). But it turns out to be questionable whether people intuit that Gettier subjects lack knowledge (Star-

mans and Friedman 2012; Turri 2012b; Weinberg, Stich and Nichols 2001).² And this should be unsurprising in light of the empirical literature on expertise, which shows that experts are especially bad at predicting what novices will do (Hinds 1999). If professional theorists of knowledge are experts in the areas of knowledge and knowledge ascription, then they'll face predictable obstacles in predicting what most people will think or say about knowledge.³

Also absent from the literature is any experimental evidence that prior explanations of lottery judgments identify psychologically relevant factors. And make no mistake about it: philosophers have explicitly said that they're trying to explain "why we typically judge" the way we do in lottery cases (DeRose 1996: 569), and that they're proposing accounts of "the relevant psychological forces driving the relevant" judgments (Hawthorne 2004: 14; see also Vogel 1990: section IV). Many of their proposals generate testable predictions. We test them. The present paper, therefore, is just as much a matter of *armchair psychology meeting experimental philosophy* as it is of *armchair philosophy meeting experimental psychology*.

We will consider three previous proposals offered to explain skeptical lottery judgment.

- *The justification account*: unjustified belief inhibits knowledge ascription. That is, if people think that you're unjustified in

² See Buckwalter 2012 for other examples of experimental work calling into question conventional wisdom about what's supposedly obvious to anyone competent with the concept of knowledge

³ See Buckwalter unpublished ms. for a discussion of professional philosophical intuitions in light of the literature on expertise.

thinking that P,⁴ they will deny that you know P. In basic lottery cases, people think your belief that the ticket will lose is unjustified, so they deny that you know it will lose. (Nelkin 2000 and Sutton 2007: 48–53; compare Williamson 2000)

- *The chance account*: chance of error inhibits knowledge ascription. That is, if people think that there is a chance that you're wrong about P, then they will deny that you know P. In basic lottery cases, people think that there's a chance that you're wrong about the ticket losing, so they deny that you know it will lose. (Cohen 1988: 196; Lewis 1996: 557)
- *The statistical account*: unanchored statistical inference inhibits knowledge ascription. That is, if people recognize that you believe P based on statistical grounds unanchored by relevant observation, then they will deny that you know P. In basic lottery cases, people recognize that you believe that the ticket will lose based on unanchored statistical grounds, so they deny that you know it will lose. (Harman 1968; compare Nelkin 2000: 396ff)⁵

We will also propose and test a new account of our own.

- *The formulaic account*: formulaic expression inhibits knowledge ascription in basic lottery cases.

Here is the plan for the paper. Section 2 reports an experiment

⁴ We use 'P' and 'Q' as placeholders for declarative sentences or that-clauses.

⁵ Cohen (1988: 106) attributes skeptical judgment in lottery cases to "the statistical nature of our reasons." But it turns out that, on Cohen's view, this is just a mechanism for making the chance of error salient, and it is the chance of error that really explains skeptical judgment. Writes Cohen, "When the chance of error is salient, we are reluctant to attribute knowledge. Statistical reasons of the sort [possessed] in the lottery case make the chance of error salient."

that tests whether people share the skeptical judgment in basic lottery cases. It also provides an initial test of the justification and chance accounts of skeptical judgment. Section 3 reports an experiment that tests whether people share the nonskeptical judgment in testimonial lottery cases. It also provides a more pointed test of the chance account. Section 4 reports an experiment that tests the statistical account. Section 5 reports an experiment that provides an initial test of the formulaic account. Section 6 reports an experiment that further tests the formulaic account. Section 7 is a general discussion of the significance of our findings.

2. Experiment 1: skeptical judgment in basic lottery cases and the justification account's demise

We begin by reporting a simple experiment designed to test two things. First, it tests whether participants tend to deny knowledge in basic lottery cases. Second, it tests whether either the justification account or the chance account can help to explain skeptical judgment in lottery cases.

Participants (N=45, 69% males) were recruited and tested using an online platform (Qualtrics and Amazon Mechanical Turk) and compensated \$.25 for approximately two minutes of their time. Participants were 29 years old on average.⁶ Participants were located in the United States and 91% listed English as a native language. Participants were not allowed to retake the survey and re-

⁶ An omnibus analysis of variance (ANOVA) revealed no effect of sex or age. The same is true of all other experiments reported in this paper.

peat participation was prevented by screening Mechanical Turk Worker IDs. Participants read a simple story and then answered a series of comprehension and test questions, followed by a brief demographic questionnaire. Test and comprehension questions were always asked in the same order; the order of response options was rotated randomly. Participants who failed comprehension questions were excluded from the analysis. We followed these same procedures in all the studies reported in this paper.

Participants read this story:

Lois is checking out at the grocery store. The clerk says to her, "Do you want to buy a lottery ticket?" Lois answers, "No thanks — I'm not going to buy a losing lottery ticket." And Lois is right: the ticket is a loser.

Participants answered these dichotomous test questions pertaining to knowledge, justification, and the chance of error:

Lois _____ that the ticket is a loser. [knows/only believes]

Lois is _____ in believing that the ticket is a loser.
[justified/unjustified]

Was there at least some chance, no matter how small, that the ticket was a winner? [yes/no]

Upon answering the dichotomous knowledge question and the dichotomous justification question, respectively, participants were asked to rate how confident they were in their answer. Responses were collected on a 1–10 scale, anchored with “not at all confident” (=1) and “completely confident” (=10).

Answers to the dichotomous knowledge and justification questions were scored either +1 (knows, justified) or -1 (only believes, unjustified). In each case, we combined the answer to the dichotomous question with the confidence rating by multiplying them. The result is a *weighted knowledge score* and a *weighted justification score*, each of which fell on a 20-point scale ranging from -10 (maximum knowledge- or justification denial) to +10 (maximum knowledge- or justification ascription).

The results demonstrate that the skeptical lottery judgment is widely shared in basic lottery cases. The vast majority (91%) of participants judged that Lois only believes that the ticket is a loser. This is significantly more than could be expected by chance.⁷ The mean weighted knowledge score (-7.36) was significantly below midpoint on the scale (which, again, ranges from -10 through +10).⁸

In light of that resounding result, let's examine participant response to the justification question. This will tell us whether the justification account might be on the right track. The justification account says that in basic lottery cases people deny knowledge because they think that 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. That is, almost all participants should answer "no" to the justification question and the mean weighted justification score should be very low. The results were highly unfavorable to the justification account. A very strong major-

⁷ Binomial test, $p < .000001$.

⁸ One-sample t-test, $t(44) = -8.9$, $p < .000001$.

ity (80%) answered “yes” to the justification question, which is more than could be expected by chance.⁹ The mean weighted justification score (+6.13) was much higher than the mean weighted knowledge score and, importantly, significantly higher than midpoint on the scale.¹⁰ These results strongly suggest that when people deny knowledge in lottery cases, it’s not because they think justification is absent.

Next let’s examine participant response to the chance question. This will tell us whether the chance account might be on the right track. The chance account says that in basic lottery cases people deny knowledge because they think that there’s a chance you’re wrong. So if the chance account is correct, then we should expect almost all participants to answer that there was a chance that the ticket was a winner. That is, almost all participants should answer “yes” to the chance question. The results were highly favorable to the chance account. The vast majority (96%) of participants answered “yes” to the chance question, which is what we would expect if the chance account explains skeptical lottery judgment.

In summary, our findings from Experiment 1 suggest three things. First, the skeptical lottery judgment is widely shared in basic lottery cases. Second, the justification account of the skeptical lottery judgment is almost certainly false. Third, the chance account does a good job of explaining the skeptical judgment in this case and so merits further consideration.

⁹ Binomial, $p < .0001$.

¹⁰ One-sample t-test, $t(44) = 5.97$, $p < .000001$.

3. *Experiment 2: nonskeptical judgment in testimonial lottery cases and the chance account's demise*

In this section we report an experiment designed to do three things. First, it seeks to replicate our finding from Experiment 1 on *skeptical* judgment in basic lottery cases. Second, it tests whether people tend to judge lottery cases differently when the protagonist concludes that the ticket lost based on testimony (as opposed to the long odds). That is, it tests whether people do tend toward *nonskeptical* judgment in testimonial lottery cases. Third, it further tests the chance account of skeptical lottery judgment.

Participants (N=143, 51% males) were assigned to one of three conditions: Odds, News, and Odd News. Participants were 32 years old on average. Ninety-six percent listed English as a native language. Participants in Odds read a basic lottery case where the ticket owner doesn't watch a newscast but simply bases her belief on the long odds. Participants in News read a story about a ticket owner who watches the evening newscast of the winning lottery numbers; no odds or chances are ever mentioned. Participants in Odd News read a similar story, except that this time the ticket owner recalls the odds of a newscaster misreporting the winning number and bases her belief on that. All participants answered comprehension questions and two test questions, a knowledge question and a chance question, analogous to the questions from Experiment 1. Participants also rated how confident they were about their answer to the knowledge question.

Here are the three stories (manipulations underlined):

[NEWS]¹¹ Ellen bought a ticket in this week's Super Lotto. Her numbers are 49-20-3-15-37-29-8. Ellen just finished watching the evening news and they reported that a completely different number won. It was the same newscaster that reports the winning number every week on the local channel that Ellen watches. On that basis, Ellen concludes that her ticket lost. And she is right: her ticket lost.

[ODD NEWS] Ellen bought a ticket in this week's Super Lotto. Her numbers are 49-20-3-15-37-29-8. Ellen just finished watching the evening news and they reported that a completely different number won. And she recalls from her statistics class that there is only a 1-in-10,000,000 (one-in-ten-million) chance that a newscaster will misreport the winning number. On that basis, Ellen concludes that her ticket lost. And she is right: her ticket lost.

[ODDS] Ellen bought a ticket in this week's Super Lotto. Her numbers are 49-20-3-15-37-29-8. Ellen wasn't able to watch the evening news where they reported which number won. But she recalls from her statistics class that there is only a 1-in-10,000,000 (one-in-ten-million) chance that a Super Lotto ticket will win. On that basis, Ellen concludes that her ticket lost. And she is right: her ticket lost.

¹¹ We use caps to name narrative elements, and we often name narrative elements after the experimental conditions they were used in. This eases exposition and helps readers keep track of which stories appeared in which conditions, while avoiding confusion between the experimental conditions and the stories. (Participants never saw the labels.)

Turning now to analyzing the results, there was an overall effect of condition on knowledge ascription (see Table 1).¹² Next we'll look for three things. First, we'll check whether the pattern of knowledge-ascription in Odds replicates the pattern observed in Experiment 1. Second, we'll check whether knowledge ascription in News displays the predicted nonskeptical pattern. Third, we'll check whether the chance account correctly predicts the overall relationship between knowledge ascription and response to the chance question, paying special attention to the results in Odd News.

First, did the pattern of knowledge ascription in Odds replicate the pattern observed in Experiment 1? Yes, it did. As before, a strong majority (80%) denied that Ellen knows that the ticket lost, which falls significantly above chance.¹³ And the mean weighted knowledge ascription (-5.7) fell significantly below midpoint.¹⁴

Second, did knowledge ascription in News display the predicted nonskeptical pattern? It did so beautifully. In News we see the mirror image of Odds. A strong majority in News (80%) answered that Ellen knows that the ticket lost, which is significantly above chance.¹⁵ Mean weighted knowledge ascription was also significantly above the midpoint (+5.78).¹⁶

¹² For the dichotomous knowledge question: $X^2(df=2, N=143)=33.74$, $p<.000001$. For the weighted knowledge ascription: ANOVA, $F(2)=22.87$, $p<.000001$.

¹³ Binomial, $p<.001$.

¹⁴ One-sample t-test, $t(39)=5.03$, $p<.0001$. We should also acknowledge that dichotomous responses were lower than in Experiment 1 (Binomial, $p=.048$), though mean weighted scores did not differ from scores in that experiment (one-sample t-test, $t(39)=1.465$, $p=.151$).

¹⁵ Binomial, $p<.001$.

¹⁶ One-sample t-test, $t(40)=4.812$, $p<.0001$.

Third, did the chance account correctly predict the overall relationship between knowledge ascription and response to the chance question? The chance account says that people deny knowledge in basic lottery cases because they think that there's a chance you're wrong about the ticket losing. The chance account fits well with the results from the Odds condition: few people ascribed knowledge, and most people affirmed the chance of error. The chance account fits less well with the results from News: most people ascribed knowledge, while a middling percentage affirmed the chance of error. Most importantly, the chance account fits very poorly with the results from Odd News. In Odd News, a majority answered "yes" to the chance question and a majority also ascribed knowledge, *both* at rates significantly higher than expected by chance.¹⁷ The mean weighted knowledge ascription is also significantly above the midpoint.¹⁸ This is hard to reconcile with the chance account's claim that people deny knowledge because they think that there's a chance that the protagonist is wrong.

Further difficulties for the chance account arise when comparing responses across conditions. Chance judgments in Odd News and Odds don't differ significantly,¹⁹ whereas knowledge ascription in the two conditions does differ significantly.²⁰ Moreover, knowledge ascription in Odd News and Odds doesn't differ significantly

¹⁷ Binomial, both $ps \leq .015$.

¹⁸ One-sample t-test, $t(61)=2.49$, $p=.016$.

¹⁹ Fisher's exact test, $p=.748$.

²⁰ For the dichotomous knowledge question: Fisher's, $p<.00001$. For the weighted knowledge ascription: ANOVA: $F(1)=26.4$, $p=.000001$, $hp^2=.21$.

(although by one measure the difference is trending),²¹ whereas chance judgments in the two conditions do differ significantly.²² If the chance account were on the right track, we shouldn't observe these outcomes.

| | Odds | Odd News | News |
|-------------------------------|------|----------|-------|
| Mean weighted knowledge score | -5.7 | +2.73 | +5.78 |
| Ascribing knowledge (%yes) | 20% | 66% | 80% |
| Chance of error (%yes) | 88% | 90% | 39% |

Table 1: Experiment 2: Comparison across conditions of the percentage of participants answering “yes” to the knowledge and chance questions, the mean weighted knowledge scores (derived by multiplying dichotomous knowledge choice by confidence), and percentage of participants affirming chance of error.

To put these difficulties another way, the chance account predicts that the rate at which participants ascribe knowledge should be inversely proportional to the rate at which they affirm a chance of error. Odds and News roughly fit this pattern, but Odd News doesn't.

We anticipate the following objection to our criticism of the chance account. Arguably a fairer test of the chance account would begin by eliminating from the analysis all participants who denied that there was a chance that Ellen's ticket won. For, it could be argued, those participants rejected a basic premise of the story by rejecting the error possibility.²³ With those participants eliminated,

²¹ For the dichotomous knowledge question: Fisher's, $p=.124$. For the weighted knowledge ascription: ANOVA: $F(1)=3.37$, $p=.07$, $hp^2=.032$. We note that the p -value on the ANOVA is trending and probably would turn significant with a larger sample size. Proponents of the chance account might take some comfort in this.

²² Fisher's, $p<.000001$.

²³ There are other ways to interpret such a denial. For example, these parti-

the chance account predicts that the remaining participants (who all affirmed a chance of error) should overwhelmingly deny knowledge.

However, following through on this suggestion wreaks greater havoc on the chance account (Table 2). For after we eliminate participants who answered “no” to the chance question, the rate of knowledge ascription in Odd News and News is identical. Moreover, if we combine the remaining participants in Odd News and News for purposes of analysis, the aggregate rate of knowledge ascription (63%) is significantly above chance,²⁴ and the aggregate mean weighted knowledge score (+2.26) is significantly above midpoint.²⁵ But now that we’re analyzing *only* participants who *affirm* the possibility of error, the chance account can’t explain this enormous disparity with these results and those observed in Odds.

| | Odds | Odd News | News |
|-------------------------------|-------|----------|-------|
| Mean weighted knowledge score | -6.86 | +2.09 | +2.44 |
| Ascribing knowledge (%yes) | 14% | 63% | 63% |
| N= | 35 | 56 | 16 |

Table 2: Experiment 2: Including only participants who answered “yes” to the chance question. Comparison across conditions of mean weighted knowledge scores and the percentage of participants answering “yes” to the dichotomous knowledge question.

In summary, our findings from Experiment 2 taught us three things. First, we replicated the skeptical result from Experiment 1,

participants might be interpreting “chance” as “genuine chance” or “meaningful chance” or “chance that should be taken into account” for planning purposes. We won’t pursue the matter here.

²⁴ Binomial, $p=.044$.

²⁵ One-sample t-test, $t(71)=2.06$, $p=.043$.

again observing that a very strong majority share the *skeptical* judgment in basic lottery cases. Second, a very strong majority share the *nonskeptical* judgment in basic testimonial lottery cases. This pair of results confirms that philosophers have mainly gotten the role of epistemology of lotteries correct. Third, the chance account of skeptical lottery judgments faces some problems. Of course, it's consistent with our findings that the chance account captures a small part of what explains skeptical lottery judgment. We don't claim to have ruled that out. Neither do we rule out more sophisticated versions of the chance account or more sophisticated ways of testing its viability. Nevertheless, our findings in this experiment motivate us to seek alternatives.

4. *Experiment 3: skeptical judgment in other statistical cases and the statistical account's demise*

In this section we evaluate the statistical account of skeptical lottery judgment in light of the results from Experiment 2. We then report an experiment designed to further test the statistical account.

It might initially seem that Experiment 2 also provides evidence against the statistical account of skeptical judgment in basic lottery cases. For Ellen's statistical inference in Odd News very closely resembles her statistical inference in Odds. In each case she recalled that there was a 1-in-10,000,000 chance of error and on that basis concluded that the ticket lost. Despite Ellen's use of statistical reasoning in each case, people judge them *very* differently. They judge nonskeptically in Odd News but skeptically in Odds. The statistical

account can't explain why.

Tempting as that line of criticism might be, it misconstrues the statistical account. The statistical account doesn't identify the relevant factor as statistical inference *per se*. Rather, it identifies *unanchored* statistical inference as the culprit. Let us explain.

Unanchored statistical inference occurs when the relationship between the premises and the conclusion is merely statistical and *not explanatory*. Inference in testimonial lottery cases arguably involves explanation. Gilbert Harman notes that our "natural non-philosophical" view of true testimonial belief involves two assumptions (Harman 1968: 166–7). First, you believe the truth because an informant told you. Second, your informant believes what he says and "believes as he does because he" has first-hand knowledge (or was told by someone else who does have first-hand knowledge). In *Odd News* these explanatory assumptions inform Ellen's statistical inference, or so it is natural to think. The fact that she also relies on statistics doesn't obscure the explanatory anchoring. By contrast, in the statistical reasoning featured in basic lottery cases, "no explanation is involved" (Harman 1968: 167). In *Odds*, Ellen's ticket doesn't lose *because* it has only a 1-in-10,000,000 chance of winning. Nor does the ticket have a 1-in-10,000,000 chance of winning because it loses. Nor is it natural to think that Ellen believes such explanatory connections are in place.

Let's understand "observation" broadly to include both perception and consumption of testimony, and let's put the essential point this way: Ellen's conclusion is partly based on a relevant explana-

tory observation in Odd News but not in Odds. The relevant observation is the newscast, which leads us to suppose that there is a “causal or explanatory” connection between Ellen’s belief and “the fact that makes it true” (Nelkin 2000: 398). Thus the statistical account can explain the results from Experiment 2.

The statistical account’s explanation of the results from Experiment 2 is ingenious. Moreover, although its proponents haven’t touted this fact, it coheres with a well documented general tendency in human judgment whereby “causes trump statistics” (Kahneman 2011: ch. 16). Decades of experimental research show that causal information drives human judgment in ways that purely statistical information doesn’t. People typically underappreciate and often completely neglect *statistical* base rates when evaluating specific cases. By contrast, people are better at appreciating *causal* base rates and treat them as information relevant to evaluating specific cases. Even meagre causal cues tend to exert more influence than ample statistical evidence does (Ajzen 1977: 307). In Daniel Kahneman’s memorable phrase, “A mind that is hungry for causal stories finds nothing to chew on” in statistics about categories. Addicted to causation, averse to statistics — that’s the fate of intuitive human judgment.

The statistical account generates testable predictions. One prediction is that in lottery cases where the subject clearly bases her conclusion on a relevant explanatory observation, participants will tend to ascribe knowledge to her.²⁶ In a word and vividly: *if you feed the causal monster, it will come*. And it will chase our inner

²⁶ Assuming, of course, that the belief is also true and, perhaps, justified.

statistical dullard away. We tested this prediction with the following experiment.

Participants (N=133, 66% males) were 29 years old on average. Ninety-six percent listed English as a native language. Participants were randomly assigned to one of three conditions: State Odds, Mafia, and State News. Participants in State Odds read another basic lottery case, this time about the State lottery, in which Ellen performs an unanchored statistical inference based on the 1-in-10,000,000 chance of winning. Participants in State News read a testimonial lottery case in which Ellen bases her belief on the newscast. Mafia is the crucial condition because Ellen bases the conclusion on the observation that the local mafia rigged the lottery such that her ticket has only a 1-in-10,000,000 chance of winning. Although her conclusion is based on statistical inference (the long odds), it's also anchored in the causal-explanatory evidence that the lottery is rigged by the mafia. In short, Mafia involves *anchored* statistical reasoning, just like Odd News from Experiment 2 did.

Here are the stories:

[STATE ODDS] Ellen bought a ticket in this week's State Lottery. She wasn't able to watch the evening news where they reported which number won. But she is a professional statistician and correctly calculates that there is only a 1-in-10,000,000 (one-in-ten-million) chance that her ticket will win. On that basis, Ellen concludes that her ticket lost. And she is right: her ticket lost.

[MAFIA] Ellen bought a ticket in this week's State Lottery.

She wasn't able to watch the evening news where they reported which number won. But she does watch a special report that reveals that the State Lottery is rigged by members of the local mafia, so that there is only a 1-in-10,000,000 (one-in-ten-million) chance that anyone not in the mafia will win. On that basis, Ellen concludes that her ticket lost. And she is right: her ticket lost.

[STATE NEWS] Ellen bought a ticket in this week's State Lottery. She just finished watching the evening news where they reported which number won. It was the same newscaster that reports about the lottery every week, and the number announced as the winner was a completely different number than Ellen's. On that basis, Ellen concludes that her ticket lost. And she is right: her ticket lost.

Participants were then asked a series of comprehension and test questions similar to those in our previous studies. For present purposes, only the knowledge question is relevant. (We will briefly mention one point about the justification question later in the section.)

This experiment provides a good test of the statistical account because we can compare whether participant response in Mafia (anchored statistical inference) more closely resembles that of State Odds (unanchored statistical inference) or that of State News (anchored observation). Participants in Mafia were fed a juicy chunk of causal flesh — *a rigged lottery!* — whereas participants in State Odds weren't. So if participants ascribe knowledge more in Mafia

than in State Odds, then it will support the statistical account. By contrast, if there is no difference between Mafia and State Odds, or if participants decline to ascribe knowledge in Mafia, then it will undermine the statistical account. At the same time, we should expect State News to elicit the highest rate of knowledge ascription of all three conditions.

The results undermined the statistical account. There was an overall effect of condition on knowledge ascription (see Table 3).²⁷ As expected, knowledge ascription was highest in State News (89%, +7.53), far exceeding what could be expected by chance.²⁸ However, knowledge ascription in Mafia (14%, -6.5) didn't differ significantly from State Odds (27%, -4.11).²⁹ Indeed, knowledge ascription in Mafia was actually *lower* than in State Odds, and it was well below what could be expected by chance.³⁰ Given that Ellen anchors her statistical inference on an explanatorily relevant observation in Mafia, the statistical account of skeptical judgment can't explain this result.

²⁷ For the dichotomous knowledge question: $X^2(df=2, N=133)=58.37$, $p<.000001$, Cramer's $V=.662$. For the weighted knowledge score: ANOVA: $F(2)=52.96$, $p<.000001$, $hp^2=.449$.

²⁸ For the dichotomous question: binomial, $p<.000001$. For the weighted knowledge score: $t(44)=8.97$, $p<.000001$.

²⁹ For the dichotomous question: Fisher's, $p=.186$. For the weighted knowledge score: ANOVA: $F(1)=2.25$, $p=.137$.

³⁰ For the dichotomous question: binomial, $p<.000001$. For the weighted knowledge score: one-sample t-test, $t(43)=6.86$, $p<.000001$.

| | State Odds | Mafia | State News |
|--------------------------------|------------|-------|------------|
| Mean weighted knowledge score | -4.11 | -6.5 | +7.53 |
| Ascribing knowledge (%yes) | 27% | 14% | 89% |
| Ascribing justification (%yes) | 98% | 77% | 98% |

Table 3: Experiment 3: Comparison across conditions of the percentage of participants answering “yes” to the knowledge and chance questions, mean weighted knowledge scores (derived by multiplying dichotomous knowledge choice by confidence), and of percentage of participants ascribing justification.

It’s worth briefly noting that response to the justification question in State Odds replicates the main finding from Experiment 1 that doomed the justification account. While knowledge ascription in State Odds was very low, a full 98% of participants thought that Ellen’s belief was nevertheless justified. Indeed, rates of justification ascription in State News and State Odds are identical, even though rates of knowledge ascription in the two conditions differ dramatically. In Mafia too, rates of justification ascription were very high even though knowledge ascription was very low.

In summary, our findings from Experiment 3 cast serious doubt on the statistical account. Proponents of the statistical account might propose a version of the explanatory requirement that avoids these problems and withstands empirical scrutiny; we would welcome such a development. And further work might reveal greater nuance in how people attribute causal relevance to factors and assimilate that information when assessing specific outcomes. Such work could inspire more successful versions of the statistical account; again, we would welcome this. Until then, the statistical account is more a promissory note than a predictive theory, and we’re inclined to look elsewhere for an explanation of skeptical lottery

judgment.

5. *Experiment 4: relenting skeptical judgment in non-stereotypical cases and the formulaic account's promise*

In this section we propose an alternative explanation for skeptical lottery judgments. We don't propose that it *entirely* explains the rate of knowledge denial in basic lottery cases. But we submit that it's probably part of the explanation.

We suspected that there is something formulaic and stereotypical about denying knowledge in basic lottery cases. Formulaic expressions are characterized by stereotyped intonation and rhythm, familiarity, predictability, and unreflective automaticity (Van Lancker-Sidtis and Rallon 2004). Advertising campaigns by gaming boards feature formulaic slogans like 'Hey, you never know' (Hawthorne 2004: 8). And in our experience people's verbal response to basic lottery cases often comes across as clichéd. This motivated us to hypothesize that although people deny knowledge in lottery cases, they should be more likely to ascribe knowledge in similar scenarios where the protagonist's conclusion does not relate to lotteries.

To test this prediction we conducted a simple experiment. Participants (N=242, 56% males) were 33 years old on average. Ninety-seven percent listed English as a native language. Participants were randomly assigned to one of two conditions: Lotto and Phone. The story for each condition featured two people, Abigail and Stan, dis-

cussing the serial number on a ten-dollar bill. Stan specifies that the serial number is extremely likely to *not* be identical to a certain other number. In response, Abigail flat-out denies that the serial number is identical to the other number. Crucially, the two stories differ in *what the other number is*. In the story for Lotto, the other number is the winning lottery number; in the story for Phone, it is Barack Obama's mobile phone number. Participants answered comprehension and test questions analogous to those in our earlier studies. Here is the story (variations underlined and separated by a slash):

[LOTTO/PHONE] Abigail is talking with her neighbor, Stan, who is a statistician. Stan hands Abigail a bill and says, "Here is the ten dollars I owe you." Abigail looks at the bill and sees that its serial number is 5-0-6-7-4-1-6-9-8-2. Stan continues, "I made an interesting calculation. If you played that serial number in this week's lottery/dialed that serial number on a telephone, it's 99.999999% certain to lose/not be Barack Obama's mobile phone number." Abigail answers, "That serial number will not win this week's lottery/is not Obama's phone number." And Abigail was exactly right: it was a losing number/wasn't Obama's number.

The principal difference between LOTTO and PHONE is the content of Abigail's conclusion. In the one case, she concludes that the serial number isn't the winning lottery number; in the other, she concludes that it isn't Barack Obama's phone number. The chance of her being right is exactly the same in both cases. If the

formulaic account is correct, then participants will ascribe knowledge significantly more in Phone than in Lotto, because Phone is not a lottery case and so shouldn't trigger the formulaic response. This was our prediction for the experiment.

The prediction was true. There was an overall effect of condition on knowledge ascription in the predicted direction for both the dichotomous question³¹ and the weighted knowledge score (Table 4).³² On each measure, it was significantly higher in Phone and also surpassed what could be expected by chance.³³ Interestingly, although LOTTO featured unanchored statistical reasoning about losing the lottery, the rate of knowledge ascription in Lotto more than doubled from previous studies *to over 50%*. If we take as our baseline comparison the ~20% rate of knowledge ascription observed in earlier basic lottery cases involving statistical reasoning (such as Odds from Experiment 2), then this increase is statistically significant and very surprising.³⁴ The mean weighted knowledge score in Lotto was above midpoint, though not significantly. We observe a similar outcome in Experiment 6 and discuss possible explanations for it toward the end of section 6.

³¹ Fisher's, $p=.035$, one-tailed.

³² ANOVA, $F(1)=3.973$, $p=.047$.

³³ For the dichotomous question: binomial, $p=.006$. For the weighted knowledge score: one-sample t-test, $t(120)=2.937$, $p=.004$.

³⁴ Binomial, $p<.00001$.

| | Odds | Lotto | Phone | Odd News |
|-------------------------------|------|-------|-------|----------|
| Mean weighted knowledge score | -5.7 | .083 | +2.38 | +2.73 |
| Ascribing knowledge (% yes) | 20% | 50% | 63% | 66% |
| Chance of error (%yes) | 88% | 88% | 79% | 90% |

Table 4: Experiments 4: Comparison across conditions of mean weighted knowledge scores, the percentage of participants answering “yes” to the dichotomous knowledge question, and the percentage of participants affirming chance of error. Lotto and Phone conditions are from Experiment 4; Odds and Odd News conditions are from Experiment 2 and listed here for comparison.

This experiment provided some initial support for the formulaic account. The next section tests it further.

6. *Experiment 5: relenting skeptical judgment in qualitatively comparative cases*

One question about Experiment 4 is whether the results were driven by framing the probability as “99.999999% certain to lose.” A further test of the formulaic account would be to present subjects with similar cases where the probability is framed differently. For example, it could be framed qualitatively in comparison to some non-lottery-related outcome. This section reports an experiment that follows up on this.

Participants (N=200, 58% male) were 29 years old on average. Ninety-seven percent listed English as a native language. Participants were randomly assigned to one of two conditions: Comparative Lotto and Comparative Phone. The story for each condition featured the same two people, Abigail and Stan, discussing the serial number on a ten-dollar bill. Stan specifies that the serial number is “just as likely to be Brad Pitt’s mobile phone number as it is to win

this week's lottery." The probability is specified qualitatively and comparatively. In response, Abigail flat-out denies that the serial number is identical to one of those two possibilities. Crucially, the two stories differ in *which specific possibility Abigail denies* and *which she ignores* in her response. In the story for Comparative Lotto, Abigail denies that the number is the winning lottery number; in the story for Comparative Phone, she denies that the number is Brad Pitt's phone number. Participants answered comprehension and test questions similar to those in our earlier studies. For present purposes, response to only the knowledge question is relevant. Here is the story (variations underlined and separated by a slash):

[COMPARATIVE LOTTO/PHONE] Abigail is talking with her neighbor, Stan, who is a statistician. Stan hands Abigail a bill and says, "Here is the ten dollars I owe you." Abigail looks at the bill and sees that its serial number is 5-0-6-7-4-1-6-9-8-2. Stan continues, "I made an interesting calculation. That serial number is just as likely to be Brad Pitt's mobile phone number as it is to win this week's lottery." Abigail answers, "That combination will not win this week's lottery/is not Brad Pitt's mobile number." And Abigail was exactly right: that combination was a loser/it was not Brad Pitt's number.

If the formulaic account is on the right track, then rate of knowledge ascription will be significantly higher in Comparative Phone. This was our prediction about the results. By contrast, if Comparat-

ive Phone and Comparative Lotto don't differ, then it will undermine the formulaic account.

The prediction was true. There was an effect of condition on knowledge ascription in the predicted direction for both the dichotomous question³⁵ and the weighted knowledge score (Table 5).³⁶ On each measure, it was significantly higher in Comparative Phone. Moreover, once again we observed a significant difference between Comparative Lotto and the ~20% rate of knowledge ascription observed in earlier basic lottery cases involving statistical reasoning. Rate of knowledge ascription in Comparative Lotto was significantly higher than in Odds, for both the dichotomous question³⁷ and the weighted score.³⁸

| | Odds | Comp. Lotto | Comp. Phone |
|-------------------------------|------|-------------|-------------|
| Mean weighted knowledge score | -5.7 | -2.62 | +35 |
| Ascribing knowledge (% yes) | 20% | 35% | 49% |
| Chance of error (&yes) | 88% | 97% | 78% |

Table 5: Experiment 5: Comparison across conditions of mean weighted knowledge scores, the percentage of participants answering "yes" to the dichotomous knowledge question, and the percentage of participants affirming chance of error. Comp. Lotto and Comp. Phone conditions are from Experiment 5; Odds is from Experiment 2 and listed here for comparison.

These results provide further support for the formulaic account. They also demonstrate that the formulaic account's explanatory potential isn't limited to cases where the probability is explicitly and negatively framed.

³⁵ Fisher's, $p=.041$, one-tailed, Cramer's $V=.133$.

³⁶ ANOVA, $F(1)=4.078$, $p=.045$, $hp^2=.02$.

³⁷ Binomial, $p=.004$, test proportion=.2.

³⁸ One-sample t-test, $t(98)=3.56$, $p=.001$, test value=-5.7.

Experiments 4 and 5 raise at least one unanswered question: why was knowledge ascription in Lotto and Comparative Lotto significantly higher than in earlier basic lottery cases involving statistical inference, such as Odds (Experiment 2)? Explicit, negative, quantitative framing of the odds can't be the entire explanation because Comparative Lotto omits such framing. One hypothesis is that Lotto and Comparative Lotto are *presented non-stereotypically*, whereas Odds follows the stereotypical lottery "script." But this hypothesis is too coarse to fit all the data. For example, it doesn't fit with the extremely low rate of knowledge ascription observed in Mafia (Experiment 3), which seriously violates the stereotypical lottery script. Rigged lotteries are not stereotypical.

Another hypothesis attributes the higher rates in Lotto and Comparative Lotto to a much more specific aspect of the stories, namely, *the losing number's non-stereotypical source*. In Lotto and Comparative Lotto, the source is a ten-dollar bill's serial number, whereas the source in Odds is an actual lottery ticket. Importantly, the source in Mafia is an actual lottery ticket, so the low rate of knowledge ascription in Mafia doesn't threaten this hypothesis. Future research could test this hypothesis by simply matching Lotto or Comparative Lotto with a case that differs in only one respect: make the source a lottery ticket. If rates of knowledge ascription drop significantly when the source is a ticket, then it will support the hypothesis. And if rates don't drop significantly, then the hypothesis is undermined. But even if this particular hypothesis doesn't explain the difference in question, explaining the difference remains part of

fully understanding the psychology of skeptical lottery judgment.

Further research could also test other explanations for the higher rates of knowledge ascription in Phone and Comparative Phone when compared to Lotto and Comparative Lotto, respectively. We have suggested that the difference is due to formulaic response or habituation in the lottery cases. An alternative explanation for the difference is that people consider successful guesses *more plausible* in Lotto than in Phone, leading them to be more reluctant in Lotto to ascribe knowledge that the number is a loser (compare Teigen and Keren 2003). Successful guessing might be viewed as plausible in Lotto because people are *familiar* with actual instances where winning lotto numbers are guessed successfully, and because successful guessing might be viewed as the *point* or *purpose* of lotteries. But this alternative explanation is undermined by the extremely low rates of knowledge ascription in Mafia. It's not similarly plausible that Ellen wins a *rigged lottery* (when she's not the one who rigged it). Nor are people familiar with actual instances where someone wins a rigged lottery that the riggers didn't intend her to win.

In any event, let us reiterate that we don't think the formulaic account explains the entire difference between skeptical judgment in basic lottery cases and nonskeptical judgment in testimonial cases. Instead, we propose that formulaic expression accounts for part of the difference.³⁹

³⁹ We propose several additional factors in forthcoming work.

7. *General discussion*

We have shown that people share the skeptical judgment in lottery cases involving statistical reasoning, and that they also share the nonskeptical judgment in lottery cases involving testimony. In this regard, people's judgments are consistent with philosophical theorizing and philosophers have gotten the folk epistemology of lotteries mostly correct. We also tested three existing accounts for these judgments, but our findings did not support them. Contrary to the justification account, people viewed protagonists as having justified beliefs in lottery cases involving statistical reasoning. Contrary to the chance account, people viewed a protagonist as knowledgeable in a lottery case involving testimony (Odd News) even while admitting there was a chance the protagonist could have been wrong. And contrary to the statistical account, people denied that a protagonist has knowledge even when the protagonist's belief was based on anchored statistical inference (Mafia).

Our findings regarding the chance account are relevant to several lines of research in theoretical epistemology. The chance account of skeptical lottery judgment is motivated by a more general account of the nature of knowledge. Infallibilists, relevant alternative theorists, and contextualists all view knowledge as, roughly, a cognitive state that rules out chance of error. Infallibilists say that knowledge rules out *any* chance of error whatsoever, whereas relevant alternative theorists and contextualists say that knowledge rules out any *relevant* or *contextually salient* chance of error. But the results from Odd News suggest that this isn't the ordinary view

of knowledge: most participants ascribed knowledge even while admitting the chance of error. And many participants in other experiments did the same (e.g. Lotto, Phone, Comparative Lotto and Comparative Phone).

Our results regarding the justification account should be taken into consideration when evaluating the increasingly popular “knowledge-first” approach in epistemology. The knowledge-first approach tries to explain important epistemic concepts, such as evidence or epistemic probability, in terms of knowledge, thereby inverting the more traditional approach that tries to explain knowledge in terms of those other, supposedly more basic epistemic concepts (Williamson 2000). 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. For the vast majority of participants in basic lottery cases ascribe justification but deny knowledge. Of course, if knowledge-first epistemology is intended as a *prescription*, rather than as a *description* of our actual concept or practice, then our results don’t necessarily undermine it.

Our results reveal a further line of research on variations of the traditional “justified true belief” theory of knowledge. Recent empirical work suggests that the ordinary concept of knowledge is, roughly, justified true belief based on “authentic evidence” (Starmans and Friedman 2012). Authentic evidence is evidence genuinely informative about reality. If the ordinary concept of knowl-

edge is authentically justified true belief, then why don't participants ascribe knowledge in basic lottery cases, given that they acknowledge that the protagonist's belief is both true and justified? The presumptive explanation is that merely probabilistic evidence isn't viewed as genuinely informative about reality. But this is complicated by the results from Odd News. For Odd News also features probabilistic evidence even though it elicits high rates of knowledge ascription. It is also complicated by the results from Mafia. For Mafia features a lottery rigged against the protagonist's winning, which arguably is genuinely informative about whether the protagonist will win. In these ways, the "authentically justified true belief" theory of knowledge (K=AJT_B) faces challenges similar to those faced by the statistical account. Importantly, all of this points to a potential *fifth factor* in the ordinary concept of knowledge, beyond belief, truth, justification and evidential authenticity.

Finally, we also provided initial evidence for a new explanation for why some people deny knowledge in basic lottery cases: the formulaic account. Although further tests are needed to support the account, the main findings were that many people ascribed knowledge in non-stereotypically presented lottery cases and that even more people ascribed knowledge in Phone and Comparative Phone, both lottery-like cases. We also found that the formulaic account makes accurate predictions across contexts where the relevant probabilities are framed differently. This includes contexts where the probabilities are framed (i) quantitatively and explicitly, and (ii) qualitatively and comparatively.

Regardless of whether the formulaic account is supported in future experiments, this finding is of broader import. The finding suggests that theorists may want to be cautious in proposing any *general* explanation of why knowledge is not possessed in cases where the protagonist concludes, on purely statistical grounds, that a certain outcome obtains. For the findings from Phone and Comparative Phone show that many people are willing to ascribe knowledge in at least some cases matching this description.⁴⁰

Word count: 8201 (with references: 8526)

References

- Ajzen, I. (1977). Intuitive theories of events and the effects of base-rate information on prediction. *Journal of Personality and Social Psychology*, 35(5), 303–314. doi:10.1037//0022-3514.35.5.303
- Beebe, J. R., & Buckwalter, W. (2010). The Epistemic Side-Effect Effect. *Mind & Language*, 25(4), 474–498.
- Buckwalter, W. (2012). Non-Traditional Factors in Judgments about Knowledge. *Philosophy compass*, 7(4), 278–289. doi:10.1111/j.1747-9991.2011.00466.x
- Buckwalter, W. (Unpublished ms.) Expert intuition fail. University of Waterloo.
- Cohen, S. (1988). How to be a Fallibilist. *Philosophical Perspectives*, 2, 91–123.
- DeRose, K. (1996). Knowledge, assertion and lotteries. *Austral-*

⁴⁰ For helpful feedback and discussion, we thank Peter Blouw, Wesley Buckwalter, Charles Millar, Merreck Levene, and Angelo Turri. This research was supported by the Social Sciences and Humanities Research Council of Canada and an Ontario Early Researcher Award.

Asian Journal of Philosophy, 74(4), 596–605.

Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, 23(6), 121–123.

Harman, G. (1968). Knowledge, inference, and explanation. *American Philosophical Quarterly*, 5(3), 164–173.

Hinds, P. J. (1999). The curse of expertise: the effects of expertise and debiasing methods on prediction of novice performance. *Journal of Experimental Psychology: Applied*, 5 (2), 205–221.

Kahneman, D. (2011). *Thinking, Fast and Slow*. Doubleday.

Kyburg, H.E, Jr. 1961. *Probability and the logic of rational belief*. Wesleyan UP.

Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.s

Nelkin, D. (2000). The Lottery Paradox , Knowledge , and Rationality. *Philosophical Review*, 109(3), 373–409.

Myers-Schulz, B & Schwitzgebel, E. (Forthcoming.) Knowing that P without Believing that P. *Nous*.

Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. *Cognition*, 124(3), 272–83.

Teigen, K., & Keren, G. (2003). Surprises: low probabilities or high contrasts? *Cognition*, 87, 55–71. doi:10.1016/S0

Turri, J. (2012a). In Gettier's wake. In S. Hetherington (Ed.), *Epistemology: the key thinkers* (pp. 214–229). London: Continuum.

Turri, J. (2012b). Is knowledge justified true belief? *Synthese*, 184(3), 247–259. doi:10.1007/s11229-010-9773-8.

Van Lancker-Sidtis, D., & Rallon, G. (2004). Tracking the incidence of formulaic expressions in everyday speech: methods for classification and verification. *Language & Communication*, 24(3), 207–240. doi:10.1016/j.langcom.2004.02.003

Vogel, J. (1990). Are there counterexamples to the closure principle? In M. D. Roth & G. Ross (Eds.), *Doubting* (pp. 13–27).

Dordrecht: Kluwer Academic Publishers.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.