Probability Distributions of Common Repeated Events Are Misestimated

Oleg Urminsky

University of Chicago, Booth School of Business

First Version: February 2014
Current Version: December 2017

*** Please contact the author for an updated version before citing or distributing***

Correspondence concerning this article should be addressed to Oleg Urminsky, University of Chicago Booth School of Business, Chicago, IL 60637 USA.

Contact: oleg.urminsky@chicagobooth.edu

Financial support from the James M. Kilts Center for Marketing and from the True North Communications Inc Faculty Research Fund at the University of Chicago Booth School of Business is gratefully acknowledged.

Abstract

Accurately estimating the prospective probability distribution arising from repeated events with known probabilities, such as the number of heads in ten coin flips, represents a simple aptitude necessary for explicit Bayesian updating and useful in optimal decisions in the face of future uncertainty. Across elicitation methods and decision scenarios, people express beliefs that are systematically biased relative to the actual distribution. Participant beliefs reflect a “wizard-hat” shaped distribution, over-estimating the tails and under-estimating the shoulders of the distribution, relative to the actual bell-curve shape. While experts are relatively more accurate than novices, both show significant bias. The bias is not explained by regression to the mean, random error or participant heterogeneity, and is exacerbated by increasing the number of repeated events. The findings caution against assuming Bayesian belief formation in models of statistical reasoning about explicit prospective beliefs based on repeated events with known probabilities.

Keywords: Bayesian models, heuristics, inference, judgment, risk and uncertainty, statistical reasoning.

1. Introduction

Statistical reasoning provides the foundation for making future predictions and managing risk. A large literature (e.g., Tversky & Kahneman, 1974) has argued that people's statistical reasoning is often biased by using heuristics. In particular, beliefs about non-linear processes (e.g., mathematical judgment, Van Dooren, De Bock, Janssens & Verschaffel, 2008; exponential growth, Wagenaar & Sagaria, 1975; compounding interest, Stango & Zinman, 2009) are subject to linearity biases.

In contrast, an emerging literature in cognitive psychology argues that decision-making is well-represented by Bayesian models. In this view, people's decision processes are based on near-optimally representing and combining probability distribution information (Griffiths & Tenenbaum, 2006), perhaps spontaneously at the neural level (Knill & Pouget, 2004; Ma, Beck, Latham, & Pouget, 2006). Systematic errors in higher-order cognition tasks are attributed to uncertainty about the inputs or to the computational cost of highly complex probabilistic inferences (Lieder, Griffiths & Goodman, 2012; Pouget, Beck, Ma & Latham, 2013), rather than limitations in the basic component processes of simple probabilistic reasoning. Recent papers (Jones & Love, 2011; Griffiths, Chater, Norris, & Pouget, 2012; Bowers & Davis, 2012) have debated the degree to which Bayesian processes can plausibly explain performance in complex decision tasks.

In contrast, this paper directly investigates the kind of simple prospective statistical inferences that are a prerequisite for more complex Bayesian reasoning. Three experiments test how well people generate the probability distribution of known-probability repeated events, which has been characterized as a fundamental building block in statistical reasoning (e.g., the Plinko game; Goodman & Tenenbaum, 2015; Vul, 2010) and should be simple for a Bayesian decision-maker.

Across elicitation methods and decision scenarios, both novices and experts estimate a biased “wizard-hat” shaped subjective binomial distribution. They over-estimated the tails and under-estimated the shoulders of the distribution relative to the “gendarme-hat” shaped (Edgeworth’s characterization, per Stigler 1999) normative distribution.

The bias in estimates is broadly consistent with prior findings documenting errors people make in generating “sampling distributions” (Kahneman & Tversky, 1972; Peterson, DuCharme and Edwards 1968; Wheeler and Beach 1968). However, unlike prior findings, these results cannot be explained by high error variance and regression to the mean, heterogeneity, misunderstanding of general distribution characteristics, or elicitation-specific biases. The findings challenge the emerging view that the human brain is adept at optimal statistical processing for sufficiently simple tasks, with broad implications for people’s ability to accurately assess and manage risk.

2. Study 1

2.1 Method

Adult participants were recruited from Amazon MTurk to complete an online survey. A target of 900 participants were requested, yielding 867 completed surveys. Records with duplicate IP addresses, or who failed a basic attention check were removed prior to analysis, yielding 821 valid completes. The participants constituted a novice population, with less than half holding a Bachelor’s degree and only 5% identifying themselves as knowledgeable in statistics.

Participants read a hypothetical scenario, either about one of three equal-probability events which would be repeated for 10 independent trials (coin-flips, survey sampling or soccer kicks), or about the distribution of height (a control task that could be estimated based on memory rather than probabilistic inferences). To test the accuracy of beliefs about the full probability distribution

resulting from the repeated binary outcome, participants were asked to make estimates using one of three different elicitation methods, and were paid a linear accuracy incentive of up to \$1.

In two coin-flipping scenarios, participants read that they would flip a coin 10 times, winning \$1 each time that it came up heads, and getting \$0 whenever it came up tails. Participants either estimated the probability of earning each of the possible amounts (from \$0 to \$10), or made a frequentist estimate of the number of people out of 100 earning each amount. In a separate pre-test, most participants correctly estimated the 50% chance of a single coin flip coming up heads or tails based on the same cover story and elicitation method (45/49, or 92%), confirming comprehensibility of the scenario.

In the survey-sampling scenario, participants estimated the likelihood of different numbers of Coke drinkers among 10 people randomly selected from a population with equal numbers of Coke and Pepsi drinkers. The soccer scenario involved estimating the number of goals for 10 kicks of a soccer ball into a goal, from a distance that gives the kicker a 50% chance of a goal on each kick. As an additional control, all participants who were assigned an equal-probability scenario then also estimated the proportion of adult men in the U.S. in each of 11 height ranges.

Participants were randomly assigned to one of three different elicitation methods: editing an adjustable histogram (with 11 bars, one for each outcome, starting with a uniform distribution; see Orhun & Urminsky 2013), entering the probabilities as 11 numeric values that were required to sum to 100, or choosing among 6 predefined histograms (see Figure 3), with the order of response options counterbalanced. Full stimuli are provided in the online appendix. Thus, the full study used a 5 (first task: coin flipping probability, coin flipping frequentist, survey, soccer, or height) x 3 (elicitation: adjustable histogram, numeric entry, or histogram choice) design.

2.2 Results

2.2.1 Accuracy of beliefs for elicited distributions.

Combining the open-ended elicitation methods (adjustable histogram and numeric) in the three test scenarios, participants' estimates diverged substantially from the true binomial distribution (Figure 1, left panel). Participants significantly overestimated both the tails (0, 1, 2, 8, 9 and 10) and peak (5) of the distribution, and underestimated the shoulders (3, 4, 6 and 7) of the distribution (all $ps < .001$). As a result, the estimates reflected a “wizard-hat” distribution shape, rather than the correct and familiar bell-curve-like binomial, inconsistent with over-representation of extreme events (Lieder, Hsu and Griffiths 2017). While the mean of the average estimated distribution was not significantly biased (5.04 vs. 5.00, $t(418)=0.72$, $p=.47$, $d=.04$), the estimates reflected significantly higher variance (4.94 vs. 2.5, $t(418)=14.4$, $p<.001$, $d=0.70$) and more kurtosis (.09 vs. -.20, $t(399)=3.51$, $p<.001$, $d=0.18$) than the true binomial distribution.

Figure 1: Participant estimates (dashed line) exhibit the “wizard hat” shape, significantly diverging from the true binomial probabilities (histogram) in Study 1 for all participants (left), as well as for the subset of participants who made estimates that monotonically decreased from 5 (right). Error bars represent 95% confidence intervals.

This bias is not due to the elicitation procedure, as the estimated height distributions were closer to the correct bell-curve shape. Error for the height estimates were significantly lower than for the binomial distribution scenario estimates (RMSE=22.4 vs. 29.6 within-subjects, $t(418)=7.22$, $p<.001$, $d=0.35$, Figure 2). This was driven primarily by differences in the average variance of the distributions (4.21 vs. 4.94, $t(418)=4.54$, $p<.001$, $d=0.22$) as well as the average kurtosis (.29 vs. .09, $t(399)=2.04$, $p=.042$, $d=0.10$).

Figure 2: Binomial distribution estimates (black dashed line) significantly diverged from the true binomial probabilities (black solid line) more than height estimates (grey dashed line) diverged from the distribution of height (grey solid line). Error bars represent 95% confidence intervals.

2.2.2 Testing potential rationalizations.

These findings suggest a systematic bias in participants' beliefs about how individual binary probabilities aggregate, rather than general mistakes in estimating distributions. However, incorrect average estimates could arise from a mix of accurate and highly inaccurate estimates, or from accurate estimates distorted by high-variance errors (Erev, Wallsten, & Budescu, 1994).

Contrary to this interpretation, the bias was widespread. Some participants (44.9%) gave estimates that were clearly the wrong shape (15.5% of participants near-uniform, 4.5% put all the mass on one outcome, and 24.8% gave other distributions that did not monotonically decrease on both sides from a maximum at 5). Eliminating these distributions yielded 231 responses, 55.1% of the sample. Even these higher-quality estimated distributions, pre-screened for the plausibility of their shape, significantly diverged from true binomial values, overestimating the tails and the peak of the distribution, and underestimating the shoulders (all $ps < .01$, Figure 1, right panel; Average variance = 3.55 vs. 2.5, $t(230)=9.2$, $p < .001$, $d=0.60$; Average kurtosis = .42 vs. -.20, $t(230)=7.0$, $p < .001$, $d=0.46$).

In a k-means cluster analysis with one of the cluster centers constrained to the true binomial values, the near-accurate cluster was small, accounting for less than 9% of participants. All the larger clusters displayed different variations of the same basic pattern, over-estimating the tails and under-estimating the shoulders of the distribution. These analyses demonstrate that near-accurate estimates were rare, and that the findings are not an artifact of diluting a large set of accurate estimates with implausible estimates.

Alternatively, persistent over-estimation of the low probabilities in the tails could occur due to regression to the mean, particularly because of truncation of noisy but unbiased estimates at zero probability. However, such an “accurate plus error” account would not explain the fact that the

shoulders of the distribution were underestimated relative to the most likely outcome (5). To more formally test for this possibility, log-normal error models with a common error variance and with separate error variances for each true probability level were estimated. Inconsistent with the “accurate plus error” interpretation, average estimates from simulations of unbiased beliefs subject to the estimated error variances were significantly different from participants’ actual estimates (Appendix A).

2.2.3 Accuracy when choosing between distributions.

To test whether the bias extended to recognition tasks, a total of 239 participants were instead shown charts representing six distributions and asked to choose which was the most accurate, second most accurate and least accurate. The majority of participants (74%) did not identify the most probable distribution (Chart D in Figure 3), although it was chosen significantly more often than expected by chance (25.5% vs. 16.7%, $Z=3.14$, $p=.002$). In contrast, the same participants were more likely to correctly identify the most accurate distribution among the same six distribution in the control height task (39.3% vs. 25.5%, $\chi^2=12.6$, $p<.001$). People also failed to identify the second most accurate histogram (Graph F, 11.3%) and the least accurate histogram (Graph B, 8.4%) for the binomial tasks, both chosen significantly below chance levels ($Z=2.38$, $p=.02$ and $Z=4.63$, $p<.001$, respectively).

The average of the binomial distributions selected as most accurate significantly overestimated the probability of the tails, and significantly underestimated the probability of the shoulders, although the most likely outcome (5) was unbiased. The average variance was significantly higher than in the true distribution (4.21 vs. 2.5, $t=12.12$, $p<.001$, $d=0.78$). The average kurtosis, however, was lower than in the true distribution (-0.33 vs. -0.2, $t=8.15$, $p<.001$, $d=0.53$), unlike for elicited distributions.

Figure 3: Percent of participants identifying each chart as the most accurate or as the least accurate in Study 1. Most participants did not identify Chart D as the most accurate (only 26%) or Chart B as the least accurate (9%). In order from most to least accurate: Charts D, F, A, E, C, B.

2.2.4 Robustness across scenarios and elicitations.

A regression analysis (Table 1) tested whether differences across scenarios and elicitation tasks affected the accuracy of the estimated distribution. Error (RMSE) was significantly lower in the recognition task (choosing a chart), than in the two production methods (generating a histogram or providing numerical estimates). Participant accuracy in the coin-flip scenario did not vary

depending on whether people estimated probabilities or frequencies ($\beta=-0.07$, $p=.97$), contrary to the view that frequentist framing increases accuracy (Gigerenzer, 1991; Cosmides & Tooby, 1996).

Table 1

Regression analysis predicting error (RMSE) in participant estimates based on task characteristics (N=658). The baseline categories for the regression were choosing among the six graphs in the coin probability task.

	β	SE	T	p
Constant	10.01	1.51	6.67	.000
<i>Elicitation Method:</i>				
Adjustable Histogram	17.81	1.60	11.14	.000
Numeric Entry	18.97	1.58	12.05	.000
<i>Scenario:</i>				
Coin: Probability vs. frequentist	-0.07	1.78	0.04	.969
Soccer	2.20	1.78	1.23	.219
Survey	2.67	1.87	1.43	.153

The observed over-estimation of the distribution's tails is consistent with over-estimating the likelihood of long sequences of identical outcomes. While such positive-recency beliefs have been identified in skill-based contexts, the opposite negative recency beliefs have typically been found instead in luck-based contexts (Oskarsson, Van Boven, McClelland & Hastie, 2009). However, accuracy did not differ between the coin-flip (luck-based) and soccer (skill-based) scenarios ($\beta=2.20$, $p=.22$), suggesting that recency beliefs in general, and the representativeness heuristic in particular (Tversky & Kahneman, 1971), were not significantly incorporated into people's predictions of the distributions.

These results can be contrasted with prior research on inaccuracy of estimated distributions. In the course of studying conservatism in Bayesian probability revision (i.e., inferences about which distribution samples had come from, Edwards 1967), researchers elicited prior distributions. These distributions, usually for unequal probabilities, were characterized as "too flat," over-

estimating the tails and underestimating the center of the distribution (Peterson, DuCharme & Edwards 1968; Wheeler & Beach 1968; see Benjamin, Rabin and Raymond 2015 for a related model).

This result has been explained as an artifact of the elicitation technique (Teigen 1974a), and therefore generalized to other estimated distributions, such as height (Teigen 1974b, Vlek 1973). Furthermore, these findings could have also been due to people avoiding low probabilities in general (DuCharme 1970), making noisy estimates leading to regressive averages (Costello & Watts 2014), or a subset of people misunderstanding the task and making flat (equal probability) estimates.

In contrast, the results of Study 1, using the simplest odds (50-50), finds inaccurate distributions that are not merely “too flat,” as in the prior literature. Instead, participants overestimated the tails specifically relative to the shoulders (but not the center) of the distribution. Furthermore, these results involve specifically probabilistic reasoning, as they did so for the binomial probability distribution significantly more than when estimating height distributions. Because these findings are robust to elicitation method and persist when eliminating flat estimates and controlling for heterogeneity and error variance, Study 1 provides novel evidence of a systematic bias in beliefs about prospective probability distributions that cannot be explained as an artifact of elicitation.

3. Study 2

The next study tested whether education, particularly formal training in statistics, enables unbiased estimates.

3.1 Method

Two versions of a paper-and-pencil survey presented the single-player coin flip scenario from Study 1 to an expert sample and a novice sample. Experts were recruited at the 2012 Society for Judgment and Decision Making conference, to complete a survey in exchange for a large Toblerone candy bar. Recruitment during all breaks between talks yielded a total of 96 completed surveys. Novices were recruited from a community sample, at the University of Chicago's downtown lab, to complete the same paper and pencil survey in exchange for \$2, with a target of approximately 130 participants, yielding 132 surveys.

In the survey, participants were either asked to shade in the distribution on a pre-printed histogram template with an unlabeled y-axis or to make judgments among the same six histograms as in Study 1. After excluding participants who did not complete the target task, the sample consisted of 92 experts (62 drawing, 30 histogram choice) and 127 novices (64 drawing, 63 histogram choice).

3.2 Results

In the histogram-shading version, both the novice and expert samples provided estimates that diverged significantly from the correct distribution (Figure 4), over-estimating the tails and under-estimating the shoulders and center of the distribution. While novices' and experts' mean estimates were unbiased (4.73 and 4.96 vs. 5.0), the estimated variance was higher than actual for both (novices: 5.82 vs. 2.5, $t=8.53$, $p<.001$, $d=1.07$; experts: 5.30 vs. 2.5, $t=9.8$, $p<.001$, $d=1.24$). The experts' estimated distributions, while biased, were significantly more accurate than the novices' estimated distributions (RMSE=17.0 vs. 33.3, $t=5.18$, $p<.001$, $d=0.93$).

In the histogram-choice version, both novices and experts were significantly more likely than chance to identify the correct graph (novice: 27% vs. 17%, $Z=1.86$, $p=.04$, $d=0.52$; expert 43% vs. 17%, $Z=3.78$, $p<.001$, $d=1.08$). However, in both samples the chosen distributions

overestimated the variance (novice: 4.56 and experts: 3.98 vs. 2.5, both $ps < .001$). The experts were non-significantly more likely than novices to select the correct histogram ($\chi^2 = 2.49, p = .16$).

A separate study with MBA students trained in statistics ($N = 87$) similarly found estimates in the survey-sampling scenario that were significantly biased but more accurate than the Study 1 novices (see Study A5 in the online appendix).

Figure 4: Expert participants (black dashed line) significantly diverged from the true binomial probabilities (solid line), but were more accurate than non-expert participants (grey dashed line). Error bars represent 95% confidence intervals.

4. Study 3

The next study tested the robustness of the findings to the number of repeated outcomes in the scenario, among both lay respondents and experts.

4.1 Method

A target of 180 participants were requested from Amazon Mturk. From 175 completed surveys, records with duplicate IP addresses, or who failed a basic attention check were removed prior to analysis, as well as 45 participants from a separate test (reported in the appendix), yielding 127 valid completes. In addition, 71 MBA students trained in statistics completed a one-page survey. Of these, 11 were excluded due to totals not adding to 100%, yielding 187 completes overall.

Participants read one of three versions of the coin-flip game, involving either 10 flips which could each earn a dollar (as in Study 1), 21 flips which could each earn 50 cents (online sample only), or 109 flips which could each earn 10 cents. Participants then estimated the probability distribution of 11 ranges of potential winnings (each representing 9% of the outcomes) either by estimating quantities or via the adjustable histogram. The outcomes were comparable in terms of potential winnings across the three versions, with the lowest range including all outcomes of less than \$1, and the highest range including all outcomes of \$10 or more.

4.2 Results

In all three conditions, the tails of the distribution (winning less than \$3, winning \$8 or more) were significantly over-estimated (10 flips: 29.1% estimated vs. 10.9% actual, $t(71)=7.2$, $p<.001$, $d=.85$; 21 flips: 40.1% estimated vs. 2.7% actual, $t(41)=8.9$, $p<.001$, $d=1.4$; 109 flips: 34.3% estimated vs. under .01% actual, $t(72)=12.5$, $p<.001$, $d=1.5$). Likewise, while the actual probability of the most likely outcome (winning between \$5 and under \$6) increased substantially with the number of coin flips (from 24.6%, to 33.6% to 66.2%), the estimated probability did not vary with the number of flips (24.7% vs. 20.8% vs. 22.3%). As a result, the average error was significantly higher for estimates of 21 flips than for 10 flips (RMSE = 37.2 vs. 22.5, $t(112)=4.11$,

$p < .001$, $d = .75$) and for 109 flips than for 21 flips (RMSE = 51.9 vs. 37.2, $t(113) = 3.88$, $p < .001$, $d = .71$).

Comparing the MBA experts with the Mturk lay sample (for the 10 and 109 flip conditions), the expert sample was slightly more accurate (RMSE = 33.7 vs. 39.9; in regression controlling for scenario $t(142) = 1.85$, $p = .066$, $d = .28$). However, the expert sample was biased, significantly over-estimating the tails of the distribution (10 flips: 24.9% estimated vs. 10.9% actual, $t(30) = 3.3$, $p = .002$, $d = .59$; 109 flips: 24.5% estimated vs. under .01% actual, $t(28) = 7.1$, $p < .001$, $d = 1.3$). The experts' estimated probability for the median outcome range also did not reflect higher probability for more flips (28.0% vs. 25.1%), yielding significantly higher average error for estimates of 109 flips than for 10 flips (RMSE = 46.3 vs. 21.9, $t(58) = 5.66$, $p < .001$, $d = 1.18$).

These results are very similar to the “law of small numbers” findings of Kahneman and Tversky (1972) for elicited distributions of number of boys born on a day when either 10, 100 or 1000 children were born. However, their findings were also subject to the alternative explanations discussed in Study 1. Their results could be explained by noisy estimates (Costello & Watts 2014), uncertainty about the underlying probability (Benjamin, Rabin and Raymond 2015), or heterogeneity, with some people making flat estimates, particularly if the prevalence of noisy or flat estimates increased with task complexity (e.g., with the sample size in the scenario).

In contrast, the results of Study 3 cannot be explained by these factors. The results persist excluding the non-monotonic and near-flat estimates ($N = 104$ remaining; see Appendix C). Likewise, as in Study 1, average estimates from simulations of unbiased beliefs subject to the estimated error variances were different from participants' actual estimates, and more so for the large sample sizes, inconsistent with an “accurate plus error” explanation (Appendix A). These

results suggest that the bias in beliefs about probability distributions may be exacerbated rather than improved by increasing the sample size, even among experts.

5. General Discussion

The ability to aggregate repeated events into accurate probability distributions has been proposed as a basic building block of near-optimal human cognition. These findings demonstrate a robust and problematic bias in making explicit estimates of such distributions, not explained by heterogeneity or random error. The resulting “wizard-hat” subjective distributions overestimate the least likely outcomes and underestimate moderately likely outcomes, relative to the actual binomial distribution for repeated independent events. This bias is very robust, but does vary in degree with elicitation methods, number of repeated events and level of expertise.

There are multiple ways in which the bias could occur. People may note that there is only one way to get no heads in ten flips, but that the other outcomes are increasingly more possible. They may then (incorrectly) reason that there are k times more ways to have exactly one coin come up heads, and k times more than that to get two heads, all the way up to the most likely outcome. A doubling sequence of this type ($k=2$) yields results similar to the average monotonic estimates in Figure 2. Future research should investigate this and other heuristics that could cause the bias.

Prior work has identified over-weighting of rare events when probabilities are explicit and underweighting of rare events when probabilities are learned from experience (Hertwig, Barron, Weber & Erev, 2004; Hertwig & Erev, 2009). Prospective decisions, such as betting on the outcome of a sports game, represent a common, distinct and understudied third type of probabilistic decisions. In such cases, as in these studies, people need to infer net outcome probabilities for future events from their beliefs about the component events (i.e., likelihood of

scoring), often without experience. The systematic bias identified in this paper may therefore represent an additional and distinct cause of over-weighting small probabilities, when the outcome probabilities are not known and are inferred from repeated events. This is in contrast with the popularly-accepted view that unlikely future events arising from a confluence of factors (e.g. “black swans”; Taleb, 2010), are generally under-estimated.

In fact, these findings suggest that forecasts involving similar prospective judgments (e.g., inflation forecasts; Goldstein & Rothschild, 2014; retirement preferences; Delavande & Rohwedder, 2008) may likewise over-estimate the likelihood of rare outcomes arising from low-probability independent joint events, impeding effective risk management. For example, the probabilities of never having a car accident or of getting all the numbers correct when playing the lottery are likely to be over-estimated, relative to more likely outcomes, even when people are well-calibrated about the likelihood of each component event.

The ability to combine and update probabilities into accurate distributions of outcomes underlies the normative Bayesian standard of inference and is central to people’s performance in a broad range of higher-order cognition tasks. These findings pose a challenge for those economic (eg., Erdem & Keane, 1996; Grossman, Kihlstrom & Mirman, 1977) and psychological (Dawes, 1989; McGuire & Kable, 2013; Viscusi, 1985) theories which explain behavior by assuming that people are able to make near-optimal decisions by efficiently integrating information into accurate explicit probability beliefs. The systematic bias in people’s subjective distributions of outcomes, estimated from simple repeated events with full information about the underlying probability, represents a reason for doubt that people generally do so efficiently.

References

- Benjamin, D. J., Rabin, M., & Raymond, C. (2015). A model of nonbelief in the law of large numbers. *Journal of the European Economic Association*, 14(2), 515-544
- Bowers, J. S., & Davis, C. J. (2012). Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin*, 138(3), 389.
- Cosmides, L., & Tooby, J. (1996). Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty. *Cognition*, 58(1), 1-73.
- Dawes, R. M. (1989). Statistical criteria for establishing a truly false consensus effect. *Journal of Experimental Social Psychology*, 25(1), 1-17.
- Delavande, A., & Rohwedder, S. (2008). Eliciting subjective probabilities in Internet surveys. *Public Opinion Quarterly*, 72(5), 866-891.
- Edwards, W. (1968). Conservatism in human information processing. In B. Kleinmuntz (Ed.), *Formal Representation of Human Judgment*. New York: Wiley, 17-52.
- Erdem, T., & Keane, M. P. (1996). Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Science*, 15(1), 1-20.
- Goldstein, D. G., & Rothschild, D. (2014). Lay understanding of probability distributions. *Judgment and Decision Making*, 9(1), 1-14.
- Gigerenzer, Gerd. (1991). "How to make cognitive illusions disappear: Beyond 'heuristics and biases', in W. Stroebe & M. Hewstone (Eds.) *European Review of Social Psychology* (Vol. 2, pp. 83–115), Chichester: Wiley.

- Goodman, N. D. and Tenenbaum, J. B. (electronic). *Probabilistic Models of Cognition*, Chapter 2. Retrieved April 23, 2015 from <http://probmods.org>.
- Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science*, 17(9), 767-773.
- Griffiths, T. L., Chater, N., Norris, D., & Pouget, A. (2012). How the Bayesians got their beliefs (and what those beliefs actually are): Comment on Bowers and Davis (2012). *Psychological Bulletin*, 138(3), 415-422.
- Grossman, S. J., Kihlstrom, R. E., & Mirman, L. J. (1977). A Bayesian approach to the production of information and learning by doing. *The Review of Economic Studies*, 533-547.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8), 534-539.
- Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, 13(12), 517-523.
- Jones, M., & Love, B. C. (2011). Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behavioral and Brain Sciences*, 34(04), 169-188.
- Kahneman, D., & Tversky, A. (1972) Subjective probability: A judgment of representativeness. *Cognitive Psychology* 3 (3), 430-454.
- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *TRENDS in Neurosciences*, 27(12), 712-719.
- Lieder, F., Griffiths, T., & Goodman, N. (2012). Burn-in, bias, and the rationality of anchoring. In *Advances in Neural Information Processing Systems* (pp. 2690-2798).

- Lieder, F., Griffiths, T. L. & Hsu, M. (2017). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources, *Psychological Review*, forthcoming
- Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, 9(11), 1432-1438.
- McGuire, J. T., & Kable, J. W. (2013). Rational temporal predictions can underlie apparent failures to delay gratification. *Psychological Review*, 120(2), 395.
- Orhun, A. Y., & Urmitsky, O. (2013). Conditional projection: How own evaluations influence beliefs about others whose choices are known. *Journal of Marketing Research*, 50(1), 111-124.
- Oskarsson, A. T., Van Boven, L., McClelland, G. H. & Hastie., R. (2009). What's next? Judging sequences of binary events. *Psychological Bulletin*, 135(2), 262–285.
- Peterson, C. R., DuCharme, W. M. & Edwards, W. (1968) Sampling distributions and probability revisions. *Journal of Experimental Psychology* 76 (2): 236.
- Pouget, A., Beck, J. M., Ma, W. J., & Latham, P. E. (2013). Probabilistic brains: knowns and unknowns. *Nature Neuroscience*, 16(9), 1170-1178.
- Stango, V., & Zinman, J. (2009). Exponential growth bias and household finance. *The Journal of Finance*, 64(6), 2807-2849.
- Stigler, S. M. (1999). *Statistics on the table: The history of statistical concepts and methods*. Harvard University Press, 95
- Taleb, N. N. (2010). *The Black Swan: The Impact of the Highly Improbable Fragility*. Random House Digital, Inc..
- Teigen, K. H. (1974a). Subjective sampling distributions and the additivity of estimates. *Scandinavian Journal of Psychology*, 15(1), 50-55.

- Teigen, K. H. (1974b). Overestimation of subjective probabilities. *Scandinavian Journal of Psychology*, 15(1), 56-62.
- Tversky A. & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin*, 76(2), 105-110.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124-1131.
- Van Dooren, W., De Bock, D., Janssens, D., & Verschaffel, L. (2008). The linear imperative: An inventory and conceptual analysis of students' overuse of linearity. *Journal for Research in Mathematics Education*, 311-342.
- Viscusi, W. K. (1990). Do smokers underestimate risks? *Journal of Political Economy*, 1253-1269.
- Vlek, C. A. (1973). Coherence of human judgment in a limited probabilistic environment. *Organizational Behavior and Human Performance*, 9(3), 460-481.
- Vul, E. (2010). Sampling in human cognition (Doctoral dissertation, Massachusetts Institute of Technology).
- Wagenaar, W. A., & Sagaria, S. D. (1975). Misperception of exponential growth. *Perception & Psychophysics*, 18(6), 416-422.
- Wheeler, G., & Beach, L. B. (1968) Subjective sampling distributions and conservatism. *Organizational Behavior and Human Performance* 3 (1), 36-46.

Probability Distributions of Common Repeated Events Are Misestimated -- APPENDIX

Oleg Urminsky

University of Chicago, Booth School of Business

Appendix A: Comparison of Alternative Explanations

Heterogeneity.

To investigate the role of heterogeneity in the biased estimates, a cluster analysis was conducted, with one cluster center fixed at the true distribution values. K-means cluster analyses were conducted on the eleven estimates comprising the estimated distribution for different numbers of segments, with the six segment solution providing the best fit (Calinski and Harabasz 1974). The cluster centers for the six segment solution were then supplemented with an additional cluster, setting this cluster's center equal to the true binomial distribution values. The k-means algorithm was then re-run as classification only, yielding a seven cluster solution, with one cluster representing near-accurate estimates.

The four largest segments (plotted in Figure 1) accounted for 93% of participants. The accurate-by-definition segment (Segment 4) was small, accounting for less than 9% of participants. All the other large segments displayed the same basic pattern, over-estimating the tails of the distribution and under-estimating the shoulders of the distribution. The remaining small segments (totalling 7% of the sample) were characterized by very high tail estimates.

Figure 1: Average estimates in accurate segment (S4, dark line) vs. larger segments.

This analysis demonstrates that near-accurate estimates were rare, and that the bias is consistent across individuals, rather than being an artifact of combining very different types of misestimates.

Random-error alternative explanations.

Could the results arise from accurate probabilities estimated with unbiased random errors (Costello and Watts, 2014)? Such a “regression to the mean” explanation is unlikely, given that the highest-probability outcome (e.g, 5 out of 10) was not under-estimated. However, to test this possibility more formally, two such models were defined and calibrated (full details of the estimation are provided in Appendix B).

Let $s_{ki} = \pi_k + \varepsilon_{ki}$, where π_k is the true probability of outcome k , and $\varepsilon_{ki} \sim N(0, \sigma^2)$ is the random error for person i estimating outcome k (out of the m possible outcomes). Assume that the estimate y_{ki} has the form:

$$y_{ki} = \frac{e^{s_{ki}}}{\sum_k e^{s_{ki}}}$$

This constrains the estimates to be between 0 and 1 (non-inclusive), and represents a “true plus error” homogeneous-error model, with one parameter, the variance σ^2 .

For the first model (Err1), the error variance is assumed to be the same for all participants, and for all outcomes being estimated. To estimate σ^2 , randomly assign six of the outcomes to set O1 and the other five outcomes to set O2 for each person, and calculate an index:

$$M = \sum_{k \in O1} -(1/6) \ln(y_{ki}) + \sum_{k \in O2} (1/5) \ln(y_{ki})$$

Under the model assumptions, $\text{Var}(M) \approx .3667\sigma^2$. Based on the estimated error variance, the average y_k each outcome can be simulated, and compared to the actual estimates.

In the second model (Err6), the assumption that the error variance is the same for all outcomes is loosened to symmetric errors: $\varepsilon_{ki} \sim N(0, \sigma_k^2)$ for $k \leq m/2$ and $\varepsilon_{ki} \sim N(0, \sigma_{m-k}^2)$ for $k > m/2$. Separate indices $M_{ki} = \ln(y_{ki}) - \ln(y_{hi})$ are computed for each person i and outcome $k \leq m/2$ (setting $h=10-k$). Each variance parameter σ_k^2 can then be estimated as $.5 * \text{Var}(M_{ki})$, for $k \leq m/2$. If m is even, for $k=m/2$ a different index $M_{ki} = \ln(y_{ki}) - (1/10) \sum_{j \neq k} \ln(y_{ji})$ is also computed. The error variance σ_k is estimated from $\text{Var}(M_{ki}) = \sigma_k^2 - .01 \sum_{j \neq 5} \sigma_j^2$.

For both models, the estimated error variance and the true outcome probabilities were used to simulate estimates for 10,000 participants. The averages of these simulated estimates are plotted below, along with the actual probabilities and the mean estimates from Study 1.

Figure 2: Estimated (solid gray line) vs. best fit error-model predictions (constant variance, dotted line; heterogeneous variance, dashed line) vs. actual values (solid black line) for 10 flips.

Simulated data from both the single-parameter homogenous-variance model (Err1) and the less parsimonious heterogeneous-variance model (Err6) both differ from the true values, suggesting that random error can induce bias. However, both models also diverge substantially from the participants' estimates. In particular, participants gave higher estimates for the tails (i.e. 0, 1, 9, 10) and lower estimates for the shoulders of the distribution (3, 4, 6, 7) than in the estimated error-models. These results strongly suggest that error variance alone cannot explain the results. The same analysis yielded similar results in Study 3, as illustrated in Figures C1-C3 below.

Constant-Ratio Sequential Inference

A relatively simple constant-ratio heuristic can generate the observed “wizard-hat” shaped distribution. In the 10 coin-flip task, people may assume that the number of ways of getting each outcome is double the less likely adjacent outcome. Thus, in the histogram, the middle bar for the 5-heads outcome would be estimated as roughly double the height of the 4 and 6

heads bars; those bars would in turn be roughly double the height of the 3 and 7 heads bars, and so on.

More formally, let m be the number of flips and r be the constant ratio. The count C_k of ways to get each outcome k would then be r^k left-hand side of the distribution ($k \leq m/2$) and symmetrically $r^{(m-k)}$ for the right-hand side ($k > m/2$). The estimated probability for each outcome would then be the share of possible outcomes $p_k = C_k / \sum_j C_j$.

In fact, applying this heuristic to the 10 coin-flip task using $r=2$ (doubling), yields heuristic-based estimates that are strikingly similar to the actual estimates made by the people in Study 1 who gave monotonically decreasing estimates from 5.

Figure 3: Average of monotonic estimates in Study 1 (N=231, solid gray line) vs. doubling constant ratio heuristic (dotted line) vs. actual values (solid black line).

Thus, this intuitive and easily implemented (but incorrect) heuristic could generate the observed results for monotonic estimated distributions. However, other heuristics might also be able to generate these results. Furthermore, the other participants (who gave non-monotonic distributions) are most likely using other, even cruder and less accurate, heuristics.

The constant-ratio heuristic alone would also not explain the insufficient sensitivity of estimates in Study 3 to the number of coin flips. People might have made ratio-based inferences on the number of categories tested (which was constant across the conditions with 10 or more flips), rather than the number of outcomes. Alternatively, the subjective ratio r could vary with the number of outcomes. It would be useful for future research to investigate the viability of different heuristics to explain variation in estimated distributions in scenarios varying the key parameters.

References

- Calinski, T, Harabasz J (1974) A dendrite method for cluster analysis. *Communications in Statistical Theory and Methods* 3(1):1–27
- Costello, F., & Watts, P. (2014). Surprisingly rational: Probability theory plus noise explains biases in judgment. *Psychological review*, 121(3), 463–480.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4), 297–323.

Appendix B: Details of Error Models

This appendix provides derivations for the model-based tests reported in Appendix A. The goal is to provide a baseline model of a decision maker who has unbiased (i.e. accurate) beliefs about the binomial distribution that are then reported with random error. The model takes into account the subsequent constraints imposed by the elicitation procedure: all values need to be between 0 and 100 and need to sum to 100. Two forms of the model were derived. Model 1 will assume that the pre-constraints error variance is the same for all outcomes, while Model 2 loosens that assumption, allowing the error variance to differ for outcomes with different true values.

Let $p_0 - p_{10}$ be the true binomial probabilities that are then estimated for outcomes \$0 to \$10. The estimates of each person i are written as y_{0i} to y_{10i} .

Homogeneity of Variance (one parameter) Model

For each estimated outcome k made by each person i , a latent variable is defined as $\exp(s_{ki}) = \exp(p_k + \varepsilon_{ki})$, where $\varepsilon_{ki} \sim N(0, \sigma^2)$ iid is the random error that occurs in the estimation process. Note that this assumes that there is a single error variance, not only across people, but also across estimated outcomes (i.e. for the probability of winning zero in the game or the probability of winning \$5 in the game). We will loosen this assumption in the next section.

Next, each person's estimates are assumed to be of the form $y_{ki} = 100 * \exp(s_{ki}) / \sum_k \exp(s_{ki})$. This constrains the estimates y_{ki} to be positive and to sum to 100.

In the studies reported in the paper, $k=0$ to 10. The 11 estimates are assigned either to set S1 (6 elements) or set S2 (five estimate) at random. Once this is done, an index M_i can be computed:

$$M_i = \sum_{k \text{ in } S1} -(1/6) \ln(y_{ki}) + \sum_{k \text{ in } S2} (1/5) \ln(y_{ki})$$

Thus:

$$M_i = \sum_{k \text{ in } S_1} -(1/6)s_{ki} + \sum_{k \text{ in } S_2} (1/5)s_{ki} + (6/6) \ln [\sum_k \exp(s_{ki})] - (5/5) \ln [\sum_k \exp(s_{ki})]$$

$$M_i = \sum_{k \text{ in } S_1} -(1/6) (p_k + \varepsilon_{ki}) + \sum_{k \text{ in } S_2} (1/5) (p_k + \varepsilon_{ki})$$

Defining a constant $C_i = \sum_{k \text{ in } S_1} -(1/6) p_k + \sum_{k \text{ in } S_2} (1/5) p_k$ for each person i , M_i can be rewritten as:

$$M_i = C_i + \sum_{k \text{ in } S_1} -(1/6) \varepsilon_{ki} + \sum_{k \text{ in } S_2} (1/5) \varepsilon_{ki},$$

Since the ε_{ki} are iid and $\varepsilon_{ki} \sim N(0, \sigma^2)$, $\text{Var}(M_i) = (6 (1/6)^2 + 5 (1/5)^2) \sigma^2 = .3667 \sigma^2$

Based on this result, the estimation strategy is to compute the index M_i for each person using different randomizations of estimates into S_1 and S_2 . The sample variance of the indices is then computed, and the model variance parameter is estimated as $\sigma^2 = \text{Var}(M_i)/.3667$.

Using this method, the error variance in the model for Study 1 was estimated as $\sigma^2 = 1.384$. Then, 10,000 observations were simulated using the accurate values of p_{ki} and the estimated σ^2 . These simulated results can then be compared to the empirical data, with any remaining differences in mean estimates unattributable to the one-parameter “true plus error” model.

Heterogenous Variance (six parameter) Model

Alternatively, it is possible that each separate point on the binomial distribution has its own error variance. For example, perhaps there is less error variance when estimating the probability of an even split (5 head out of 10 tosses) than either extreme outcome (10 heads, or 10 tails). Formally, the model would then be extended to include outcome-specific error variances:

$$\exp(s_{ki}) = \exp(p_k + \varepsilon_{ki}) \text{ where } \varepsilon_{ki} \sim N(0, \sigma_k^2).$$

This can be simplified somewhat by making a reasonable assumption of symmetry in the error variances of corresponding outcomes. In the eleven-outcome case, with $k = 0 \dots 10$, this would be written as: $\sigma_k^2 = \sigma_{10-k}^2$.

Using a similar strategy as the prior model, five indices are computed:

$$M_{ki} = \ln(y_{ki}) - \ln(y_{hi}), \text{ for } k = 0 \text{ to } 4, \text{ where } h = 10 - k.$$

$$M_{ki} = s_{ki} - s_{hi} - \ln [\sum_k \exp(s_{ki})] + \ln [\sum_k \exp(s_{ki})]$$

$$M_{ki} = (p_k + \varepsilon_{ki}) - (p_h + \varepsilon_{hi}) = \varepsilon_{ki} - \varepsilon_{hi}, \text{ since } p_k = p_h$$

Since $\text{Var}(\varepsilon_k) = \text{Var}(\varepsilon_h) = \sigma_k^2$, the model variance for outcome $k < 5$ can be estimated as $.5 \text{ Var}(M_{ki})$.

For $k = 5$, it will be necessary to construct the index differently:

$$M_{5i} = \ln(y_{5i}) - (1/10) \sum_{k \neq 5} \ln(y_{ki})$$

By a similar argument, this will reduce to:

$$M_{5i} = \varepsilon_{5i} - (1/10) \sum_{k \neq 5} \varepsilon_{ki}$$

Since the errors $\varepsilon_{ki} \sim N(0, \sigma_k^2)$ are assumed to be iid:

$$\text{Var}(M_{5i}) = \sigma_5^2 - .01 \sum_{k \neq 5} \sigma_k^2$$

Thus, an estimated of the error variance σ_5^2 can be computed as:

$$\text{Var}(M_{5i}) - .01 \sum_{k \neq 5} \sigma_k^2, \text{ based on the previously computed estimated } \sigma_k^2 \text{ for } k \neq 5.$$

Using this approach for the Study 1 data yielded estimates of:

$$\sigma_0^2 = \sigma_{10}^2 = 10.325, \sigma_1^2 = \sigma_9^2 = 4.442, \sigma_2^2 = \sigma_8^2 = 3.568,$$

$$\sigma_3^2 = \sigma_7^2 = 3.084, \sigma_4^2 = \sigma_6^2 = 1.842, \sigma_5^2 = 11.510.$$

Simulating 10,000 cases based on a “true + error” model using these error variance estimates yielded average estimates that significantly differed from the empirical data in Study 1, as shown in Appendix A.

Appendix C: Additional Methods Details and Analyses

Study 1

Method Details

Participants from Amazon Mechanical Turk were recruited, with the goal of having approximately 900 participants. From a total of 867 completed surveys, records with duplicate IP addresses, or who failed a basic attention check were removed (prior to any analysis), yielding 821 valid completes. After completing an unrelated survey, participants read the scenario and were asked to make the estimates.

Height data was compiled from the National Health and Nutrition Examination Surveys¹ (NHANES) conducted from 1999-2012. Data was screened for respondents who had a valid height in cm, were male and were 18 years old or older, yielding 18,628 records. These were converted to inches, and coded into the two-inch intervals used in the survey.

The study used a 5 (first task: coin flipping probability, coin flipping frequentist, survey, soccer, or height) x 3 (elicitation method: adjustable histogram, numeric entry or choice among 6 histogram graphs). In the numeric entry and adjustable histogram versions, the order of the outcomes (e.g., 0 first or 10 first) was counterbalanced. In the choice among graphs version, the order of the six graphs was rotated, yielding four different displays (i.e. counterbalancing top vs. bottom and left column vs. right column).²

¹ http://www.cdc.gov/nchs/nhanes/nhanes_questionnaires.htm

² In the graph choice version of the survey scenario, one of the four displays was not asked due to a programming error.

Additional Analyses

Table C1: *Estimates for each outcome compared to actual values, for elicited binomial outcomes.*

	Correct	All Data (N=419)			Monotonic Distributions (N=231)		
		Avg. Estimate	<i>T</i>	<i>p</i>	Avg. Estimate	<i>T</i>	<i>p</i>
Outcome=0	0.10	4.47	8.52	<.001	1.46	12.15	<.001
Outcome=1	0.98	5.32	13.21	<.001	3.12	13.11	<.001
Outcome=2	4.39	6.08	7.49	<.001	5.09	3.97	<.001
Outcome=3	11.72	7.63	19.68	<.001	8.37	15.51	<.001
Outcome=4	20.51	11.63	26.02	<.001	14.69	16.90	<.001
Outcome=5	24.61	28.37	3.39	<.001	34.88	10.75	<.001
Outcome=6	20.51	11.95	25.35	<.001	14.67	16.65	<.001
Outcome=7	11.72	8.01	17.62	<.001	8.30	16.60	<.001
Outcome=8	4.39	6.16	7.64	<.001	4.87	2.99	.003
Outcome=9	0.98	5.93	11.63	<.001	2.98	14.16	<.001
Outcome=10	0.10	4.44	9.69	<.001	1.57	12.57	<.001

Table C2: *Comparison of estimates to actual values separately in each scenario.*

	Correct	Coin flip probability (N=106)			Coin flip frequentist (N=98)		
		Avg. Estimate	<i>t</i>	<i>p</i>	Avg. Estimate	<i>T</i>	<i>p</i>
Outcome=0	0.10	4.66	5.09	<.001	5.86	4.09	<.001
Outcome=1	0.98	5.50	6.77	<.001	5.42	6.07	<.001
Outcome=2	4.39	5.94	4.51	<.001	6.00	3.96	<.001
Outcome=3	11.72	7.78	11.20	<.001	7.80	10.40	<.001
Outcome=4	20.51	11.96	14.49	<.001	11.84	11.44	<.001
Outcome=5	24.61	24.50	0.07	.948	26.86	1.16	.249
Outcome=6	20.51	12.51	13.92	<.001	11.71	12.71	<.001
Outcome=7	11.72	8.56	8.70	<.001	8.44	7.90	<.001
Outcome=8	4.39	6.59	5.41	<.001	6.44	4.38	<.001
Outcome=9	0.98	6.52	6.76	<.001	5.47	6.47	<.001
Outcome=10	0.10	5.49	5.08	<.001	4.17	6.56	<.001
Mean Outcome	5.00	5.12	1.31	.192	4.94	0.56	.575
Variance	2.50	5.55	7.95	<.001	5.24	7.91	<.001
Kurtosis	-0.20	-0.14	0.59	.558	0.09	1.64	.105
RMSE	--	26.52			28.02		

Table C2: *Comparison of estimates to actual values separately in each scenario. (Continued)*

	Correct	Survey (N=114)			Soccer (N=101)		
		Avg. Estimate	<i>t</i>	<i>p</i>	Avg. Estimate	<i>T</i>	<i>p</i>
Outcome=0	0.10	2.09	4.70	<.001	5.62	4.52	<.001
Outcome=1	0.98	4.16	8.25	<.001	6.35	6.61	<.001
Outcome=2	4.39	5.76	2.67	.009	6.66	4.49	<.001
Outcome=3	11.72	6.89	10.68	<.001	8.14	7.83	<.001
Outcome=4	20.51	11.85	11.62	<.001	10.84	15.75	<.001
Outcome=5	24.61	37.36	4.86	<.001	23.76	0.40	.694
Outcome=6	20.51	11.97	11.22	<.001	11.58	13.82	<.001
Outcome=7	11.72	6.75	11.95	<.001	8.46	7.06	<.001
Outcome=8	4.39	4.71	0.99	.322	7.08	4.38	<.001
Outcome=9	0.98	4.06	7.17	<.001	7.85	5.39	<.001
Outcome=10	0.10	4.38	4.31	<.001	3.67	4.60	<.001
Mean Outcome	5.00	5.08	1.37	.173	4.99	0.08	.939
Variance	2.50	4.29	5.65	<.001	4.75	7.76	<.001
Kurtosis	-0.20	0.27	2.53	.013	0.15	1.88	.064
RMSE	--	32.46			31.01		

Table C3: *Height distribution estimates, first tasks only.*

	Correct	Height Estimates First (N=120)		
		Estimate	<i>t</i>	<i>P</i>
Below 5' 1"	0.60	2.40	5.39	<.001
5' 1" to under 5' 3"	2.74	3.76	3.04	.003
5' 3" to under 5' 5"	8.50	6.46	5.31	<.001
5' 5" to under 5' 7"	17.53	10.78	12.88	<.001
5' 7" to under 5' 9"	24.98	16.54	12.44	<.001
5' 9" to under 5' 11"	23.58	21.69	2.09	.039
5' 11" to under 6' 1"	14.12	16.97	4.37	<.001
6' 1" to under 6' 3"	5.94	10.36	8.87	<.001
6' 3" to under 6' 5"	1.57	5.47	12.35	<.001
6' 5" to under 6' 7"	0.36	3.50	9.28	<.001
6' 7" and above	0.07	2.07	7.33	<.001

Table C4: *Within-subjects comparison of prediction error in binomial scenario and height distribution estimates (N=419).*

	Average Error			
	<i>Binomial</i>	<i>Height</i>	<i>t</i>	<i>p</i>
Outcome=0	+4.38	+2.04	4.16	<.001
Outcome=1	+4.34	+1.19	9.38	<.001
Outcome=2	+1.68	-2.36	14.14	<.001
Outcome=3	-4.09	-7.10	8.11	<.001
Outcome=4	-8.88	-7.02	2.84	<.001
Outcome=5	+3.76	-1.84	4.89	<.001
Outcome=6	-8.56	+1.59	21.29	<.001
Outcome=7	-3.71	+3.87	20.76	<.001
Outcome=8	+1.77	+4.22	8.38	<.001
Outcome=9	+4.95	+3.20	3.97	<.001
Outcome=10	+4.34	+2.21	4.54	<.001
RMSE	29.57	22.35	7.21	<.001

Table C5: *Between-subjects comparison of binomial scenario and height distribution estimates, first tasks only.*

	Binomial (N=419)		Height (N=120)		<i>t</i>	<i>p</i>
	Mean	SE Mean	Mean	SE Mean		
Mean Outcome	5.04	0.05	4.94	0.06	0.98	0.329
Variance	4.94	0.17	4.07	0.16	2.67	0.008
Kurtosis	0.09	0.08	0.22	0.08	0.84	0.402
RMSE	29.57	1.00	20.87	0.88	4.52	<.001

Table C6: *RMSE regression estimates, using only monotonic elicited distributions (N=470)*

	β	SE	<i>T</i>	<i>P</i>
Constant	10.92	.95	11.50	.000
<i>Elicitation Method:</i>				
Adjustable Histogram	7.18	1.10	6.52	.000
Numeric Entry	9.02	1.08	8.34	.000
<i>Scenario:</i>				
Coin: Probability vs. frequentist	-0.78	1.20	0.65	.517
Soccer	-0.10	1.22	0.08	.934
Survey	2.39	1.27	1.89	.060

Table C7: Chosen histograms (from set of 6, all with mean of 5).

	Correct	Recoded Graph Choice (N=239)		
		Estimate	<i>t</i>	<i>p</i>
Outcome=0	0.10	2.42	15.42	<.001
Outcome=1	0.98	3.05	13.46	<.001
Outcome=2	4.39	4.88	3.38	.001
Outcome=3	11.72	9.90	19.06	<.001
Outcome=4	20.51	17.76	10.01	<.001
Outcome=5	24.61	23.98	1.42	.156
Outcome=6	20.51	17.76	10.01	<.001
Outcome=7	11.72	9.90	19.06	<.001
Outcome=8	4.39	4.88	3.38	.001
Outcome=9	0.98	3.05	13.46	<.001
Outcome=10	0.10	2.42	15.42	<.001
Variance	2.50	4.21	12.12	<.001
Kurtosis	-0.20	-0.33	8.15	<.001
RMSE	--	11.15		

Table C8: Percent of respondents choosing each histogram as most accurate, second most accurate or least accurate. (Correct answer denoted in bold)

	Binomial Scenarios (N=239)			Height (N=282)		
	Most	Second most	Least	Most	Second most	Least
Chart A	6.3%	11.7%	4.2%	6.7%	14.5%	6.7%
Chart B	8.8%	24.7%	8.4%	8.9%	18.8%	9.9%
Chart C	18.8%	11.7%	33.1%	14.5%	13.8%	31.2%
Chart D	25.5%	25.5%	3.8%	40.8%	20.9%	3.5%
Chart E	32.6%	15.1%	27.2%	19.5%	19.9%	27.0%
Chart F	7.9%	11.3%	23.4%	9.6%	12.1%	21.6%

Study 2

Method Details

Expert sample. Participants were recruited at the 2012 Society for Judgment and Decision Making conference in the hall outside the meeting rooms, to complete a survey in exchange for a large Toblerone candy bar. A total of 96 participants handed in surveys. In the paper-and-pencil survey, along with other unrelated questions, participants were asked to either draw a histogram representing 10 coin flips, or to choose among the same six histograms as in Study 1, with the order of outcomes and graphs counterbalanced. Four of the participants did not complete the coin-flipping scenario, yielding 92 participants (62 doing drawings and 30 making choices).

Novice sample. Participants were recruited from a community sample, at the University of Chicago's downtown lab, to complete the same paper and pencil survey, in exchange for \$2, with a target of approximately 130 participants. After excluding 5 participants who did not complete the coin flip task, we had 64 participants who completed the drawing task and 63 who completed the histogram choice task, for a total of 127.

Additional Analyses.

Table C13: *Comparison of estimates to actual values separately in each scenario.*

	Correct	Novice sample (N=64)			Expert sample (N=62)		
		Avg. Estimate	<i>t</i>	<i>p</i>	Avg. Estimate	<i>T</i>	<i>P</i>
Outcome=0	0.10	6.89	4.72	<.001	3.58	5.23	<.001
Outcome=1	0.98	8.69	7.53	<.001	4.93	9.09	<.001
Outcome=2	4.39	9.05	6.38	<.001	7.31	7.58	<.001
Outcome=3	11.72	9.35	2.86	.006	10.81	2.34	.022
Outcome=4	20.51	11.95	9.25	<.001	14.14	11.97	<.001
Outcome=5	24.61	15.55	3.96	<.001	19.21	3.54	.001
Outcome=6	20.51	11.40	6.74	<.001	14.53	10.62	<.001
Outcome=7	11.72	9.61	1.36	.180	10.63	3.17	.002
Outcome=8	4.39	6.60	3.49	.001	7.29	7.70	<.001
Outcome=9	0.98	5.31	7.36	<.001	4.73	11.33	<.001
Outcome=10	0.10	5.59	3.38	0.001	2.86	7.99	<.001
Mean Outcome	5.00	4.73	1.59	0.116	4.96	0.62	.536
Variance	2.50	5.82	8.53	<.001	5.30	9.80	<.001
Kurtosis	-0.20	-0.73	6.12	<.001	-0.52	5.98	<.001
RMSE	--	33.28			16.96		

Table C14: *Between-subjects comparison of errors in novice and expert estimates.*

	Average Error			
	Novice	Expert	<i>t</i>	<i>p</i>
Outcome=0	+6.80	+3.48	2.07	.041
Outcome=1	+7.71	+3.95	3.34	.001
Outcome=2	+4.66	+2.92	2.09	.039
Outcome=3	-2.37	-0.91	1.58	.117
Outcome=4	-8.56	-6.37	2.03	.044
Outcome=5	-9.06	-5.40	1.32	.189
Outcome=6	-9.10	-5.98	2.11	.037
Outcome=7	-2.11	-1.09	0.63	.531
Outcome=8	+2.20	+2.89	0.93	.355
Outcome=9	+4.34	+3.75	0.86	.394
Outcome=10	+5.49	+2.76	1.62	.108

Table C15: *Between-subjects comparison of novice and expert estimates, summary statistics.*

	Novice (N=64)		Expert (N=62)		<i>t</i>	<i>p</i>
	Mean	SE Mean	Mean	SE Mean		
Mean Outcome	4.73	0.17	4.96	0.07	1.22	.225
Variance	5.82	0.39	5.30	0.29	1.08	.283
Kurtosis	-0.73	0.09	-0.52	0.05	-2.11	.037
RMSE	33.28	2.70	16.96	1.58	5.18	<.001

Table C16: *Between-subjects comparison of novice and expert estimates, summary statistics (using only monotonic elicited distributions).*

	Novice (N=17)		Expert (N=43)		<i>t</i>	<i>p</i>
	Mean	SE Mean	Mean	SE Mean		
Mean Outcome	5.09	0.09	5.04	0.03	0.70	.486
Variance	4.79	0.28	4.65	0.21	0.36	.722
Kurtosis	-0.28	0.19	-0.42	0.05	0.96	.340
RMSE	15.04	1.39	11.91	0.72	2.17	.034

Table C17: *Comparison of recoded chosen histograms (from set of 6, all with mean of 5) to actual values for each sample*

	Correct	Novice Sample (N=63)			Expert Sample (N=30)		
		Avg. Estimate	<i>t</i>	<i>p</i>	Avg. Estimate	<i>T</i>	<i>p</i>
Outcome=0	0.1	2.75	8.95	<.001	1.97	4.35	<.001
Outcome=1	0.98	3.37	8.26	<.001	3.00	4.10	<.001
Outcome=2	4.39	5.30	3.59	.001	4.80	1.10	.282
Outcome=3	11.72	10.10	9.21	<.001	10.10	4.91	<.001
Outcome=4	20.51	17.17	6.76	<.001	18.00	3.05	.005
Outcome=5	24.61	22.63	2.41	.019	24.27	0.38	.710
Outcome=6	20.51	17.17	6.76	<.001	18.00	3.05	.005
Outcome=7	11.72	10.10	9.21	<.001	10.10	4.91	<.001
Outcome=8	4.39	5.30	3.59	.001	4.80	1.10	.282
Outcome=9	0.98	3.37	8.26	<.001	3.00	4.10	<.001
Outcome=10	0.1	2.75	8.95	<.001	1.97	8.95	<.001
Mean Outcome	5	5.00	--	--	5.00	--	--
Variance	2.5	4.56	7.71	<.001	3.98	3.68	<.001
Kurtosis	-0.2	-0.35	4.42	<.001	-0.37	5.63	<.001
RMSE	--	11.16			9.19		

Table C18: Comparison of prediction error in recoded histogram-choice estimates

	Average Error		<i>t</i>	<i>p</i>
	<i>Novice</i>	<i>Expert</i>		
Outcome=0	+2.65	+1.87	1.50	.138
Outcome=1	+2.39	+2.02	.676	.500
Outcome=2	+0.91	+0.41	1.12	.264
Outcome=3	-1.62	-1.61	.01	.989
Outcome=4	-3.33	-2.51	.91	.368
Outcome=5	-1.97	-0.34	1.21	.228
Outcome=6	-3.33	-2.51	.91	.368
Outcome=7	-1.62	-1.62	.01	.989
Outcome=8	+0.91	+0.41	1.12	.264
Outcome=9	+2.30	+2.02	.676	.500
Outcome=10	+2.65	+1.87	1.50	.138
RMSE	11.16	9.19	1.34	.183

Table C19: Percent of respondents choosing each histogram as most accurate, second most accurate or least accurate. (Correct answer denoted in bold)

	Novice Sample (N=63)			Expert Sample (N=30)		
	Most	Second most	Least	Most	Second most	Least
Chart A	3.2%	7.9%	12.7%	10.0%	10.3%	8.7%
Chart B	9.5%	23.8%	9.5%	20.0%	10.3%	10.9%
Chart C	11.1%	19.0%	27.0%	13.3%	31.0%	23.9%
Chart D	27.0%	28.6%	6.3%	43.3%	20.7%	7.6%
Chart E	39.7%	12.7%	15.9%	13.3%	20.7%	21.7%
Chart F	9.5%	7.9%	28.6%	0.0%	6.9%	27.2%

Study 3

Method Details

Participants from Amazon Mechanical Turk were recruited, with the goal of having approximately 180 participants. From a total of 175 completed surveys, records with duplicate IP addresses, or who failed a basic attention check were removed (prior to any analysis), yielding 172 valid completes. Participants read one of four versions of the coin-flipping scenario from Study 1, with either 2, 10, 21 or 109 flips. Participants then made estimates, using either the numeric elicitation or the adjustable histogram (with outcome order counterbalanced) from Study 1. The 10, 21 and 109 flip scenarios were designed to be comparable, all elicited using 10 categories of outcomes and a similar range of monetary amounts (i.e. for the 21 flips, the eleven outcome categories ranged from \$0-\$.50, \$1 to \$1.50, up to \$10-\$10.50).

Additional Analyses – Online Sample Only

Table C9: *Comparison of estimated and actual distribution in the 10 flips condition.*

	Correct	10 Flips (N=41)		
		Avg. Estimate	<i>t</i>	<i>P</i>
Outcome=0	0.098	4.57	3.45	0.001
Outcome=1	0.977	6.68	4.15	<.001
Outcome=2	4.395	8.00	4.94	<.001
Outcome=3	11.719	10.26	-2.32	0.026
Outcome=4	20.508	13.44	-8.52	<.001
Outcome=5	24.609	22.20	-1.10	0.28
Outcome=6	20.508	13.22	-6.62	<.001
Outcome=7	11.719	8.67	-5.26	<.001
Outcome=8	4.395	5.76	3.06	0.004
Outcome=9	0.977	4.20	5.74	<.001
Outcome=10	0.098	3.01	5.41	<.001
Tails	10.9375	32.22	7.04	
Mean Outcome	5.00	4.72	-2.13	0.04
Variance	2.50	5.13	5.98	<.001
Kurtosis	-0.20	-0.11	0.51	0.615
RMSE	--	23.02

Figure C1: Estimates (grey line) vs. best fit error-model predictions (constant variance, dotted line; heterogeneous variance, dashed line) vs. actual values (solid line) for 10 flips.

Table C10: Comparison of estimated and actual distribution in the 21 flips condition.

	Correct	21 Flips (N=42)		
		Avg. Estimate	<i>t</i>	<i>P</i>
Outcome=0	0.001	9.67	3.92	<.001
Outcome=1	0.073	7.31	6.40	<.001
Outcome=2	1.256	7.69	6.63	<.001
Outcome=3	8.132	9.10	0.80	0.427
Outcome=4	23.719	11.83	-8.91	<.001
Outcome=5	33.638	20.81	-4.32	<.001
Outcome=6	23.719	10.69	-11.65	<.001
Outcome=7	8.132	7.45	-0.99	0.33
Outcome=8	1.256	5.40	6.80	<.001
Outcome=9	0.073	4.43	5.56	<.001
Outcome=10	0.001	5.62	2.60	0.013
Tails	2.6604	40.12	8.89	
Mean Outcome	5.00	4.57	-1.70	0.097
Variance	5.25	4.95	-0.66	0.514
Kurtosis	-0.10	0.84	2.00	0.052
RMSE	--	37.22

Figure C2: Estimates (grey line) vs. best fit error-model predictions (constant variance, dotted line; heterogeneous variance, dashed line) vs. actual values (solid line) for 21 flips.

Table C11: Comparison of estimated and actual distribution in the 109 flips condition.

	Correct	109 Flips (N=43)		
		Avg. Estimate	<i>t</i>	<i>P</i>
Outcome=0	0.000	10.26	4.41	<.001
Outcome=1	0.000	8.65	6.40	<.001
Outcome=2	0.000	8.18	10.44	<.001
Outcome=3	0.193	9.14	15.95	<.001
Outcome=4	16.717	12.05	-4.86	<.001
Outcome=5	66.182	20.52	-15.14	<.001
Outcome=6	16.717	10.23	-7.73	<.001
Outcome=7	0.193	7.30	14.26	<.001
Outcome=8	0.000	5.59	11.56	<.001
Outcome=9	0.000	4.54	8.82	<.001
Outcome=10	0.000	3.55	6.31	<.001
Tails	0.0001	40.77	11.18	
Mean Outcome	5.00	4.37	-3.56	0.001
Variance	27.25	5.92	-40.81	<.001
Kurtosis	-0.02	1.45	1.02	0.314
RMSE	--	55.61

Figure C3: Estimates (grey line) vs. best fit error-model predictions (constant variance, green line; heterogeneous variance, dashed line) vs. actual values (solid line) for 109 flips.

The 2 flip condition was not directly comparable to the other three conditions, and so it was analyzed separately. The tails of the distribution were substantially over-estimated. However, the overall error was much lower for the 2 flip condition than the 10 flip condition (14.88 vs. 23.02, $t(84)=2.54$, $p=.013$, $d=.56$), as well as compared to the 21 and 109 flip conditions (both $ps < .001$).

Table C12: Comparison of estimated and actual distribution in the 2 flips condition.

	Correct	2 Flips (N=45)		
		Avg. Estimate	<i>t</i>	<i>P</i>
Outcome=0	25.00	30.36	3.64	0.001
Outcome=1	50.00	41.73	-3.83	<.001
Outcome=2	25.00	27.91	2.34	0.024
Mean Outcome	1.00	0.98	1.47	0.149
Variance	0.50	0.57	3.20	0.003
Kurtosis	-1.00	-1.16	-2.77	0.008
RMSE	--	14.88

Figure C4: Estimates (dashed line) vs. actual values (solid line) for 2 flips.

The comparison of the different conditions shows that RMSE increases with the complexity (e.g., number of flips).

Figure C5: Error (RMSE) by the number of flips in the estimated distribution ($r = .93$).

Additional Analyses – Expert and Online Sample Combined

Similar results hold when constraining the data to only those estimated distributions that are monotonically decreasing on both sides from the middle and which are not near-flat (N=104).

Figure C6: Estimates (dashed lines) when constraining the data.

In all three conditions, the tails of the distribution (winning less than \$3, winning \$8 or more) were significantly over-estimated (10 flips: 16.1% estimated vs. 10.9% actual, $t(43)=3.3$, $p=.002$, $d=.49$; 21 flips: 18.6% estimated vs. 2.7% actual, $t(20)=5.5$, $p<.001$, $d=1.3$; 109 flips: 19.4% estimated vs. under .01% actual, $t(41)=10.7$, $p<.001$, $d=1.2$). Likewise, while the probability of the most likely outcome (winning between \$5 and under \$6) increased substantially with the number of coin flips (from 24.6%, to 33.6% to 66.2%), the estimated probability for the median outcome range did not vary with the number of flips (33.2% vs. 34.3% vs. 30.2%). As a result, the average error was significantly higher for estimates of 21 flips than for 10 flips (RMSE = 22.5 vs. 15.9, $t(63)=2.29$, $p=.025$, $d=.59$) and for 109 flips than for 21 flips (RMSE = 41.2 vs. 22.5, $t(61)=5.44$, $p<.001$, $d=1.21$).

Additional Study A4

Method

Forty nine online-participants were first asked to estimate the probabilities of the outcomes from a single round, and then the full 10 round coin-flip game from Study 1 using numeric elicitation and the adjustable the histogram elicitation.

Results and Discussion

The majority of participants (45/49, or 92%) correctly answered that there was a 50% chance of getting \$1 and a 50% chance of getting \$0. The other four responses indicated a 30%, 51%, 55% or 60% chance of winning \$1. Excluding participants who gave these incorrect responses, the bias in estimates for the 10 round game demonstrated in Study 1 was replicated. The tail outcomes (0, 1, 2, 8,9 and 10) were all estimated to have significantly higher probability than actual and the remaining outcomes (3 through 7) were all estimated to have significantly lower probabilities (all $ps < .05$).

Figure C7: Estimates (dashed line) vs. actual values (solid line) for 10 flip scenario, among those participants giving correct answer to the baseline probability question.

Additional Analyses.

Table C20: Comparison of estimates to actual values

	Correct	Participants passing screening question (N=45)		
		Avg. Estimate	<i>t</i>	<i>p</i>
Outcome=0	0.10	9.69	4.12	<.001
Outcome=1	0.98	8.58	5.66	<.001
Outcome=2	4.39	7.33	6.29	<.001
Outcome=3	11.72	8.62	6.51	<.001
Outcome=4	20.51	12.02	8.15	<.001
Outcome=5	24.61	18.67	2.42	.020
Outcome=6	20.51	10.87	10.33	<.001
Outcome=7	11.72	7.92	7.22	<.001
Outcome=8	4.39	6.57	4.19	<.001
Outcome=9	0.98	5.57	8.54	<.001
Outcome=10	0.10	4.17	6.99	<.001
Mean Outcome	5.00	4.55	2.55	.014
Variance	2.50	6.33	7.58	<.001
Kurtosis	-0.20	-0.26	0.22	.825
RMSE	--	29.41		

Additional Study A5

Method

All students enrolled in an MBA course were asked to complete the numeric elicitation version of the survey scenario from Study 1, as an in-class exercise at the start of a lecture on survey sampling. Eighty seven students, all of whom had completed at least one statistics course as a pre-requisite, completed the exercise. Because the paper-and-pencil version they completed did not force estimates to add up to 100, 21% of the estimates were rescaled to sum to 100 (multiplying estimate x_{ij} for person i by $100/\sum_j x_{ij}$).

Results and Discussion

The MBA student estimates (dark dashed line in Figure A11 below) significantly diverged from the binomial distribution for all the outcomes (solid line). While the mean of the distribution was estimated accurately, the students' estimated distribution had higher variance (3.68 vs. 2.5, $t=5.65$, $p<.001$, $d=0.61$) and more kurtosis (.43 vs. -.20, $t=2.89$, $p=.005$, $d=0.31$) than the actual distribution.

Figure C8: Novice estimates (light dashed line) vs. expert estimates (MBA students, dark dashed line) vs. actual values (solid line) for 10 flips.

Nevertheless, the MBA student estimates for most of the outcomes were significantly different from the estimates of the novice population in Study 1, made using the same scenario and elicitation procedure ($N=61$, light dashed line in Figure C8 above), and were closer to the actual distribution. Overall, the MBA students were significantly more accurate ($RMSE=16.0$ vs. 33.5 , $t=6.21$, $p<.001$, $d=1.04$) than the corresponding Study 1 participants.

Additional Analyses.

Table C21: Comparison of estimates to actual values separately in each scenario

	Correct	Novice sample (Study 1, N=61)			Expert sample (MBA, N=87)		
		Avg. Estimate	<i>t</i>	<i>P</i>	Avg. Estimate	<i>t</i>	<i>p</i>
Outcome=0	0.10	1.65	4.44	<.001	1.46	6.15	<.001
Outcome=1	0.98	4.64	6.02	<.001	3.09	8.57	<.001
Outcome=2	4.39	5.77	2.32	.024	5.07	2.88	.005
Outcome=3	11.72	6.55	9.00	<.001	9.28	8.32	<.001
Outcome=4	20.51	10.92	10.20	<.001	15.14	10.53	<.001
Outcome=5	24.61	39.57	4.16	<.001	31.68	4.60	<.001
Outcome=6	20.51	10.84	10.49	<.001	15.28	10.86	<.001
Outcome=7	11.72	6.49	10.16	<.001	9.36	8.13	<.001
Outcome=8	4.39	4.84	0.95	.347	5.08	2.92	.004
Outcome=9	0.98	4.60	5.45	<.001	3.09	8.50	<.001
Outcome=10	0.10	4.13	4.30	<.001	1.47	6.15	<.001
Mean Outcome	5.00	5.09	1.44	.156	5.00	1.22	.227
Variance	2.50	4.36	4.67	<.001	3.68	5.65	<.001
Kurtosis	-0.20	-0.09	0.43	.671	0.43	2.89	.005
RMSE	--	33.49			16.01		

Table C22: Between-subjects comparison of errors in novice (Study 1) and expert (MBA student) estimates.

	Average Error			
	Novice	Expert	<i>t</i>	<i>p</i>
Outcome=0	+1.55	+1.36	0.48	.633
Outcome=1	+3.66	+2.11	2.64	.009
Outcome=2	+1.37	+0.68	1.23	.223
Outcome=3	-5.17	-2.44	4.59	<.001
Outcome=4	-9.58	-5.37	4.24	<.001
Outcome=5	+14.96	+7.07	2.24	.027
Outcome=6	-9.67	-5.23	4.62	<.001
Outcome=7	-5.23	-2.36	5.21	<.001
Outcome=8	+0.45	+0.69	0.49	.625
Outcome=9	+3.63	+2.11	2.41	.017
Outcome=10	+4.03	+1.38	3.20	.002

Table C23: *Between-subjects comparison of novice and expert estimates, summary statistics.*

	Novice (N=61)		Expert (N=87)		<i>t</i>	<i>p</i>
	Mean	SE Mean	Mean	SE Mean		
Mean Outcome	5.09	0.06	5.00	0.00	1.64	.103
Variance	4.36	0.40	3.68	0.21	1.62	.108
Kurtosis	-0.09	0.26	0.43	0.22	1.52	.130
RMSE	33.49	2.73	16.01	1.38	6.21	<.001

Table C24: *Between-subjects comparison of novice and expert estimates, summary statistics (using only monotonic elicited distributions).*

	Novice (N=34)		Expert (N=81)		<i>t</i>	<i>p</i>
	Mean	SE Mean	Mean	SE Mean		
Mean Outcome	5.02	0.02	5.00	0.00	0.89	.377
Variance	3.31	0.35	3.44	0.18	-0.38	.705
Kurtosis	0.48	0.37	0.53	0.23	-0.12	.904
RMSE	24.80	2.26	15.57	1.45	3.46	.001

Appendix D: Study Stimuli

Study 1 Stimuli

Incentive instructions:

Please take your time, read closely and answer the next section very carefully. You are eligible to receive a bonus payment, via Mturk, based on the accuracy of your response. The questions you will answer have a factually correct answer. If you give the most accurate answer possible, you will get an extra \$1. If you give an answer that is not better than guessing at random, you will get \$0. If your answer is in between, you will get a partial bonus, based on your accuracy. So, for example, if your answer is halfway between completely correct and guessing at random, you would get 50 cents. *** IMPORTANT: THIS IS NOT HYPOTHETICAL. YOU CAN EARN AN ACTUAL ADDITIONAL PAYMENT BASED ON YOUR ACCURACY.***

SCENARIO 1: Coin flip – probability:

COIN FLIP GAME Imagine that you are going to play the following hypothetical game. In the game, you would flip a coin 10 times and you will show the experimenter the result each time. Each time that it comes up heads, you win \$1, and each time it comes up tails, you get nothing. In this game, any amount between \$0 and \$10 is possible. Using the graph below, indicate the probability of each of the outcomes from 10 coin flips. (0 = no chance, 50 = equally likely to happen or not happen, 100 = certain to happen) Since all the possible outcomes are listed, probabilities you give should add up to 100.

SCENARIO 2: Coin flip – frequentist:

COIN FLIP GAME Imagine that 100 people were going to play the following hypothetical game. In the game, each person would flip a coin 10 times and would show the experimenter the result each time. Each time that it comes up heads, the person wins \$1, and each time it comes up tails, the person gets nothing. In this game, for a given player, any amount between \$0 and \$10 is possible. Using the graph below, indicate how many of the 100 people would have each outcome. Since all the possible outcomes are listed, the number of people for each outcome you give should add up to the total of 100 people.

SCENARIO 3: Survey outcomes:

SURVEY OUTCOMES Imagine that there is a population in which exactly half of the people prefer Coke to Pepsi, and half prefer Pepsi to Coke, and no one is indifferent. Imagine that you conduct a survey with 10 people, and there is no sampling bias (everyone has an equal probability of completing the survey). What are the chances of each of the following possible outcomes? Using the graph below, please give a probability between 0% (never) and 100% (will always happen) for each outcome. Since all the possible outcomes are listed, probabilities you give should add up to 100.

SCENARIO 4: Soccer game:

SOCCER PRACTICE GAME Imagine that you are going to play the following hypothetical game. In the game, you would kick a soccer ball into a goal. The difficulty of the game would be adjusted for you personally. You would kick the ball from far away enough that, on any given kick, you have a 50% chance of getting the ball in (you are equally likely to make it into the goal as to not make it into the goal). You would try 10 kicks. In this game, any outcome between making no goals and making 10 goals is possible. Using the graph below, indicate the probability of each of the outcomes from making 10 tries. (0 = no chance, 50 = equally likely to happen or not happen, 100 = certain to happen) Since all the possible outcomes are listed, probabilities you give should add up to 100.

SCENARIO 5: Height estimation:

ESTIMATING HEIGHT Thinking only of men in the U.S. who are 18 years old or older, how common do you think each of the heights below is? Using the graph below, indicate the proportion of men in the U.S. 18 or older who are in each height range. Since all the possible heights are listed, the proportions you give should add up to 100.

Elicitation Mode: Numeric³

Winning \$0	<input type="text" value="0"/>
Winning \$1	<input type="text" value="0"/>
Winning \$2	<input type="text" value="0"/>
Winning \$3	<input type="text" value="0"/>
Winning \$4	<input type="text" value="0"/>
Winning \$5	<input type="text" value="0"/>
Winning \$6	<input type="text" value="0"/>
Winning \$7	<input type="text" value="0"/>
Winning \$8	<input type="text" value="0"/>
Winning \$9	<input type="text" value="0"/>
Winning \$10	<input type="text" value="0"/>
Total	<input type="text" value="0"/>

Elicitation Mode: Adjustable Histogram

		0 = no one											100 = everyone	
		0	10	20	30	40	50	60	70	80	90	100		
Winning \$10													9	
Winning \$9													9	
Winning \$8													9	
Winning \$7													9	
Winning \$6													9	
Winning \$5													9	
Winning \$4													9	
Winning \$3													9	
Winning \$2													9	
Winning \$1													9	
Winning \$0													9	
Total:													99	

³ Note: The outcome order was counterbalanced in all elicitation modes

Elicitation Mode: Choosing between histograms

Take a look at the charts below, and think about which one best represents the probability that each outcome would occur.

The height of each bar shows the probability (between 0 = no chance, 50 = equally likely to happen or not happen and 100 = certain to happen) of earning that amount of money, between \$0 and \$10.

On each of the charts below, the probabilities of the outcomes add up to a total of 100.

Chart A

Chart B

Chart C

Chart D

Chart E

Chart F

Elicitation: Numeric (Height)

Below 5' 1"	0
5' 1" to under 5' 3"	0
5' 3" to under 5' 5"	0
5' 5" to under 5' 7"	0
5' 7" to under 5' 9"	0
5' 9" to under 5' 11"	0
5' 11" to under 6' 1"	0
6' 1" to under 6' 3"	0
6' 3" to under 6' 5"	0
6' 5" to under 6' 7"	0
6' 7" and above	0
Total	0

Study 2 Stimuli

SCENARIO 1: 10 flips

COIN FLIP GAME Imagine that you are going to play the following hypothetical game. In the game, you would flip a coin 10 times and you will show the experimenter the result each time. Each time that it comes up heads, you win \$1, and each time it comes up tails, you get nothing. In this game, any amount between \$0 and \$10 is possible, in increments of \$1. Using the graph below, indicate the probability of each of the outcomes from 10 coin flips. (0 = no chance, 50 = equally likely to happen or not happen, 100 = certain to happen) Since all the possible outcomes are listed, probabilities you give should add up to 100.

- _____ Winning \$0
- _____ Winning \$1
- _____ Winning \$2
- _____ Winning \$3
- _____ Winning \$4
- _____ Winning \$5
- _____ Winning \$6
- _____ Winning \$7
- _____ Winning \$8
- _____ Winning \$9
- _____ Winning \$10

SCENARIO 2: 21 flips

COIN FLIP GAME Imagine that you are going to play the following hypothetical game. In the game, you would flip a coin 21 times and you will show the experimenter the result each time. Each time that it comes up heads, you win 50 cents (\$0.50), and each time it comes up tails, you get nothing. In this game, any amount between \$0 and \$10.50 is possible, in 50 cent increments. Using the graph below, indicate the probability of each of the outcomes from 21 coin flips. (0 = no chance, 50 = equally likely to happen or not happen, 100 = certain to happen) Since all the possible outcomes are listed, probabilities you give should add up to 100.

- _____ Winning \$0 to \$.50
- _____ Winning \$1 to \$1.50
- _____ Winning \$2 to \$2.50
- _____ Winning \$3 to \$3.50
- _____ Winning \$4 to \$4.50
- _____ Winning \$5 to \$5.50
- _____ Winning \$6 to \$6.50
- _____ Winning \$7 to \$7.50
- _____ Winning \$8 to \$8.50
- _____ Winning \$9 to \$9.50
- _____ Winning \$10 to \$10.50

SCENARIO 3: 109 flips

COIN FLIP GAME Imagine that you are going to play the following hypothetical game. In the game, you would flip a coin 109 times and you will show the experimenter the result each time. Each time that it comes up heads, you win 10 cents (\$0.10), and each time it comes up tails, you get nothing. In this game, any amount between \$0 and \$10.90 is possible, in increments of 10 cents. Using the graph below, indicate the probability of each of the outcomes from 109 coin flips. (0 = no chance, 50 = equally likely to happen or not happen, 100 = certain to happen) Since all the possible outcomes are listed, probabilities you give should add up to 100.

- _____ Winning \$0 to \$0.90
- _____ Winning \$1 to \$1.90
- _____ Winning \$2 to \$2.90
- _____ Winning \$3 to \$3.90
- _____ Winning \$4 to \$4.90
- _____ Winning \$5 to \$5.90
- _____ Winning \$6 to \$6.90
- _____ Winning \$7 to \$7.90
- _____ Winning \$8 to \$8.90
- _____ Winning \$9 to \$9.90
- _____ Winning \$10 to \$10.90

SCENARIO 4: 2 flips

COIN FLIP GAME Imagine that you are going to play the following hypothetical game. In the game, you would flip a coin 2 times and you will show the experimenter the result each time. Each time that it comes up heads, you win \$5, and each time it comes up tails, you get nothing. In this game, any amount between \$0 and \$10 is possible, in increments of \$5. Using the graph below, indicate the probability of each of the outcomes from 2 coin flips. (0 = no chance, 50 = equally likely to happen or not happen, 100 = certain to happen) Since all the possible outcomes are listed, probabilities you give should add up to 100.

- _____ Winning \$0
- _____ Winning \$5
- _____ Winning \$10

Study 3 Stimuli

SCENARIO 1: Histogram Drawing

COIN FLIP GAME

Imagine that you are going to flip a coin 10 times and you will show the game administrator the result each time. Each time that it comes up heads, you win \$1; each time it comes up tails, you get nothing.

In this game, any amount between \$0 and \$10 is possible. Using the graph below, indicate the probability of each of the outcomes from 10 coin flips by drawing a histogram. For each amount, shade in a bar such that the height of the bar represents the probability of that outcome.

SCENARIO 1: Histogram Drawing

COIN FLIP GAME

Imagine that you are going to flip a coin 10 times and you will show the game administrator the result each time. Each time that it comes up heads, you win \$1; each time it comes up tails, you get nothing.

In this game, any amount between \$0 and \$10 is possible. Looking at the graphs below, please think about which best represents the probability of each of the outcomes from 10 coin flips:

Chart A

Chart B

Chart C

Chart D

Chart E

Chart F

Chart that is the *closest* to the actual probabilities: _____

Chart that is the *second closest* to the actual probabilities: _____

Chart that is the *farthest* from the actual probabilities: _____

Study A4 Stimuli

COIN FLIP GAME Imagine that you are going to play the following hypothetical game. In the game, you would flip a coin once, and you will show the experimenter the result. If it comes up heads, you win \$1, and if it comes up tails, you get nothing. In this game, either \$0 or \$1 is possible. Using the graph below, indicate the probability of each of the outcomes from the single coin flip. (0 = no chance, 50 = equally likely to happen or not happen, 100 = certain to happen) Since all the possible outcomes are listed, probabilities you give should add up to 100.

_____ Winning \$0

_____ Winning \$1

Now, we would like to ask you about a DIFFERENT version of the coin flip game.

COIN FLIP GAME Imagine that you are going to play the following hypothetical game. In the game, you would flip a coin 10 times and you will show the experimenter the result each time. Each time that it comes up heads, you win \$1, and each time it comes up tails, you get nothing. In this game, any amount between \$0 and \$10 is possible, in increments of \$1. Using the graph below, indicate the probability of each of the outcomes from 10 coin flips.

(0 = no chance, 50 = equally likely to happen or not happen, 100 = certain to happen) Since all the possible outcomes are listed, probabilities you give should add up to 100.

_____ Winning \$0

_____ Winning \$1

_____ Winning \$2

_____ Winning \$3

_____ Winning \$4

_____ Winning \$5

_____ Winning \$6

_____ Winning \$7

_____ Winning \$8

_____ Winning \$9

_____ Winning \$10

Study A5 Stimuli

SCENARIO: Sampling (MBA classroom version)

Sampling Exercise:

Imagine that there is a population in which exactly half of the people prefer Coke to Pepsi, and half prefer Pepsi to Coke, and no one is indifferent.

We conduct a survey with 10 people, and there is no sampling bias (everyone has an equal probability of completing a survey).

What are the chances of each of the following possible outcomes? Please give a probability between 0% (never) and 100% (will always happen):

In our sample of 10 people, no one prefers Coke and all 10 prefer Pepsi.	Probability: _____
In our sample of 10 people, 1 person prefers Coke and 9 prefer Pepsi.	Probability: _____
In our sample of 10 people, 2 people prefer Coke and 8 prefer Pepsi.	Probability: _____
In our sample of 10 people, 3 people prefer Coke and 7 prefer Pepsi.	Probability: _____
In our sample of 10 people, 4 people prefer Coke and 6 prefer Pepsi.	Probability: _____
In our sample of 10 people, 5 people prefer Coke and 5 prefer Pepsi.	Probability: _____
In our sample of 10 people, 6 people prefer Coke and 4 prefer Pepsi.	Probability: _____
In our sample of 10 people, 7 people prefer Coke and 3 prefer Pepsi.	Probability: _____
In our sample of 10 people, 8 people prefer Coke and 2 prefer Pepsi.	Probability: _____
In our sample of 10 people, 9 people prefer Coke and 1 prefer Pepsi.	Probability: _____
In our sample of 10 people, all 10 prefer Coke and no one prefers Pepsi.	Probability: _____

TOTAL (should add up to 100):