Journal of Experimental Psychology:
Learning, Memory, and Cognition
2006, Vol. 32, No. 6, 1385–1402

0278-7393/06/$12.00 DOI: 10.1037/0278-7393.32.6.1385

Between Ignorance and Truth: Partition Dependence and Learning in Judgment Under Uncertainty

Kelly E. See
New York University

Craig R. Fox
University of California at Los Angeles

Yuval S. Rottenstreich
Duke University

In 3 studies, participants viewed sequences of multiattribute objects (e.g., colored shapes) appearing with varying frequencies and judged the likelihood of the attributes of those objects. Judged probabilities reflected a compromise between (a) the frequency with which each attribute appeared and (b) the *ignorance prior* probability cued by the number of distinct values that the focal attribute could take on. Thus, judged probabilities were *partition dependent*, varying with the number of events into which the state space was subjectively divided. This bias was diminished among participants more confident in what they learned, was strong and insensitive to level of confidence when ignorance priors were especially salient, and required ignorance priors to be salient only when probabilities were elicited (not during encoding).

Keywords: judgment, uncertainty, learning, partition dependence

An abundance of psychological research has explored how people judge the likelihood of uncertain events, such as a future rise in interest rates or the incumbent winning an upcoming election. Many of these studies have focused on heuristic processes that people use to evaluate the *nature* of possible events and the relative strength of evidence supporting them (Gilovich, Griffin, & Kahneman, 2002; Kahneman, Slovic, & Tversky, 1982). For instance, one may infer that the Acme Company is more likely to hire Alan than Richard because Alan more closely resembles the prototypical Acme worker. More recently, researchers have identified processes that people use to evaluate likelihood on the basis of the *number* of events that might occur (Fox & Clemen, 2005; Fox & Levav, 2004; Fox & Rottenstreich, 2003; Johnson-Laird, Legrenzi, Girotto, Legrenzi, & Caverni, 1999). For instance, one may conclude that Alan has a 20% chance of landing the Acme job because he is one of five candidates for a single position.

Kelly E. See, Department of Management and Organizations, Stern School of Business, New York University; Craig R. Fox, Anderson School of Management and Department of Psychology, University of California at Los Angeles; Yuval S. Rottenstreich, Fuqua School of Business, Duke University.

This work was partially supported by Grant SES-00-99209 from the National Science Foundation to Craig R. Fox. We thank the Fuqua School of Business and the Wharton School Risk Center for their support during various stages of this project. We thank Derek Koehler, Cade Massey, and especially Bob Clemen and Jack Soll for helpful feedback and suggestions. Finally, we thank Michael London of the University of North Carolina for early programming assistance.

Correspondence concerning this article should be addressed to Kelly E. See, Department of Management and Organizations, Stern School of Business, New York University, 44 West 4th Street, New York, NY 10012. E-mail: ksee@stern.nyu.edu

Reliance on the nature and reliance on the number of possible events are not mutually exclusive. To illustrate, consider the chances that a particular stranger identifies more with the Democratic than Republican platform. In the absence of further information, one might rely purely on the *number* of possible events: Because both party affiliations appear equally plausible the "ignorance prior" probability of 1/2 may appear to be the most appropriate response. Subsequently, as one learns more about the person in question, one might make adjustments based on the nature of events. For instance, one might rely on the representativeness heuristic, noting that "this person's features are more similar to my prototype of a Democratic voter," or the availability heuristic, reasoning "it is much easier to recall having met Democrats around here than Republicans."

The use of ignorance priors, though intuitively appealing, can lead to systematic bias. In particular, if people place any weight on ignorance priors, then probabilities will vary systematically with the way in which people subjectively partition the state space. To illustrate such *partition dependence*, consider the following scenario that we presented to 89 visitors of the Duke University student union:

> Fletcher Motor Sales is a small automobile dealership that sells two models of automobiles: coupes (2-door) and sedans (4-door). Fletcher employs four salespersons. It employs Carlos Tamayo as a coupe salesperson and Jennifer Burkhardt, Damon Jones, and Sebastian Cruz to sell sedans. Fletcher stocks sedans in the colors white, blue, black, and silver and coupes in red and green. A customer has just come to the Fletcher dealership to trade in his old car for a new one of the exact same model.

Note that any transaction at this dealership is characterized by the following two attributes: model (coupe, sedan) and salesperson (Carlos, Jennifer, Damon, Sebastian). One group of participants (n = 43) was asked "What is the probability that the customer trades in a coupe?" This question facilitates a twofold partition by model of car. A second group ($n = 46$) was asked "What is the probability that the customer buys a car from Carlos?" This question facilitates a fourfold partition by salesperson. Assuming participants believe that each salesperson does indeed sell only one model of automobile, the two queries identify the same event, but using two different attributes. Thus, if participants rely at all on ignorance priors cued by the target attribute, they should be biased toward 1/2 when asked about coupes and toward 1/4 when asked about Carlos.

The results indeed exhibited attribute-cued partition dependence. The median judged probability that a customer buys a coupe was .33, whereas the median judged probability that a customer buys a car from Carlos was .25, $p = .004$ (one-tailed) by Kruskal–Wallis test.[1] Moreover, a greater proportion of respondents answered 1/2 when asked about the coupe (42%) than when asked about Carlos (17%; $p = .01$, by Fisher's exact test), and a greater proportion of respondents answered 1/4 when asked about Carlos (48%) than when asked about the coupe (28%; $p = .04$). The same pattern of results appeared when we ran the study again using a within-participant design.[2]

Reliance on ignorance priors has been previously demonstrated for natural events for which participants must draw on their real-world knowledge (Fox & Clemen, 2005; Fox & Levav, 2004; Fox & Rottenstreich, 2003; Smithson & Segale, 2006). In the car dealership example, participants had little information on which to draw other than the number of salespeople and models of car (cf. Johnson-Laird et al., 1999). However, it is an open and important question whether people rely on attribute-cued ignorance priors, and therefore exhibit partition dependence, in situations in which they can readily observe the relative frequencies of events. Previous investigators have extended judgmental biases such as the alternative outcomes effect (Windschitl, Young, & Jenson, 2002) and subadditivity (Dougherty & Hunter, 2003; Koehler, 2000) into learning paradigms. Our primary purpose in this article is to likewise explore the extent to which people rely on attribute-cued ignorance priors (i.e., the number of potential events) versus observed relative frequencies (i.e., the nature of these events) in a controlled learning environment in which information concerning both the number and nature of events is readily available.

Perhaps the most natural hypothesis is that once people gain a modicum of knowledge about the true frequencies of events, they will discard the ignorance prior and rely exclusively on what they have learned about these frequencies. Consistent with this view is a substantial literature showing that when people can observe events, they are generally quite accurate at learning how often these events occur. Indeed, many authors have argued that people tend to automatically encode frequency information (e.g., Alba, Chromiak, Hasher, & Attig, 1980; Hasher & Zacks, 1984; Naveh-Benjamin & Jonides, 1986).

Alternatively, it may be that once people gain knowledge about true frequencies, they continue to be influenced by ignorance priors. Consistent with this view, there is evidence that judged relative frequencies tend to be biased toward salient reference points. For example, Varey, Mellers, and Birnbaum (1990) showed participants squares containing white and black dots; the relative frequency of the two colors was varied, and participants estimated the proportion of dots that were white (or black). Overestimation predominated for proportions less than 1/2, whereas underestimation predominated for proportions greater than 1/2 (for similar results, see Attneave, 1953; Erlick, 1964; Fiedler & Armbruster, 1994; Fox & Tversky, 1998, Study 2; Parducci & Wedell, 1986; Sheridan & Ferrell, 1974; Wickens, 1992. Hollands and Dyre, 2000, demonstrated a more general pattern of over- and underestimation around alternative reference points).

A secondary purpose of this article is to attempt to integrate the literature showing relative accuracy of frequency judgment with the literature showing systematic bias in proportion judgment by proposing that people rely on a compromise between their impression of relative frequency (the nature of events) and ignorance priors (the number of events). We propose further that the relative weight afforded these two sources of information is influenced by (a) a judge's confidence in his or her ability to estimate the relative frequency of events and (b) the salience of the ignorance prior. We study these factors using a learning paradigm because it provides an ideal environment for experimentally controlling information that is presented, measuring subjective confidence in what has been learned, and manipulating the salience of ignorance priors.

In three studies, we investigate partition dependence cued by attributes (like that illustrated by the foregoing car dealership example) in controlled learning environments. Our first study establishes that reliance on ignorance priors persists in a learning context and that this effect is diminished with subjective confidence in what has been observed. Our second study reveals that when ignorance priors are especially salient, they can bias likelihood judgment even for individuals who are highly confident in their memory for what they had observed. Our final study illustrates that ignorance priors can influence judged probabilities, provided they are salient at the time of judgment (i.e., retrieval), but they need not be salient when frequencies of events are learned (i.e., during encoding). We conclude with a discussion of factors that influence the salience of alternative partitions, the relationship between the present account and Bayesian models of likelihood judgment, and directions for future research. A formal characterization of the ignorance prior model (Fox & Rottenstreich, 2003) is provided in Appendix A.

Study 1: Establishing Reliance on Attribute-Cued Ignorance Priors

In Study 1 we investigate whether participants' recall of frequency information is biased toward attribute-cued ignorance priors. Participants observed a set of randomly ordered multiattribute events (restaurant receipts for various meals on various days). Our main hypothesis is that judged likelihood would be biased toward the ignorance prior suggested by the target attribute (meal of the

[1] Unless otherwise specified, because we make ex ante directional predictions, we use one-tailed test statistics throughout the article.

[2] We asked 76 undergraduates at the University of Southern California to complete an identically worded within-participant version of this scenario, counterbalancing the order in which each participant answered the two questions (coupe vs. salesperson). We found no order effects and the same mean and median pattern of results as in the between-participants version. The median-judged probability that a customer buys a coupe was .48, whereas the median-judged probability that a customer buys a car from Carlos was .28 ($p < .001$, by Wilcoxon test).

day or day of the week). The experiment provides an initial test of the extent to which people rely on attribute-cued ignorance priors versus true frequencies when they can readily observe the latter.

Second, we investigate the role of subjective knowledge and confidence on use of the ignorance prior. Previous research in a decision analysis context (Fox & Clemen, 2005) has shown a stronger bias toward ignorance priors among participants who rated themselves as less knowledgeable. Likewise, in the present investigation we expected that bias toward ignorance priors would be more pronounced in this learning context among participants who felt less confident in their memory and less knowledgeable about likelihood estimation tasks.

Finally, we explore whether reliance on ignorance priors is affected by response mode. Some authors have argued that biases in judged probability disappear when people are instead asked to judge relative frequencies (Gigerenzer, 1991, 1996; for a rebuttal, see Kahneman & Tversky, 1996), whereas others have observed that biases in judged probability persist in relative and absolute frequency judgment, though their magnitude may diminish slightly (e.g., Koehler, 2000; Rottenstreich & Tversky, 1997; Tversky & Kahneman, 1983; Tversky & Koehler, 1994).

Method

Participants. We recruited 267 undergraduates at the University of Pennsylvania to participate in a set of unrelated experiments lasting a total of approximately 1 hr in exchange for course credit.

Design and procedure. The experiment was administered on computer. Participants were randomly assigned to conditions that varied by response mode. In one condition ($n = 134$), participants were asked to assess the probability that a randomly selected object would have various features. In the other condition ($n = 133$), participants were asked to recall the proportion of events that had these features.

Participants were told that they would view information about a large number of meals that a person named Joe had eaten at a local restaurant.[3] Participants were shown an example of a receipt with two pieces of information (i.e., two attributes of the restaurant visit): the meal category (breakfast, lunch, or dinner) and the day of the week (Sunday, Monday, . . ., Saturday). Participants were then told that each of Joe's receipts would be quickly flashed on the screen, and they were instructed to try to get an intuitive sense of how often Joe ate particular meals at the restaurant and how often he ate at the restaurant on particular days of the week. They were also told that after viewing the receipts they would be asked to judge the likelihood of the events they observed. We assumed that for most participants the "meals" attribute would suggest a threefold partition and an ignorance prior of 1/3, whereas the "days of the week" attribute would suggest a sevenfold partition and an ignorance prior of 1/7.

The receipts were white rectangles with a black outline that were approximately 1.5-in. wide × 2.5-in. high (3.8-cm wide × 6.4-cm high). The background of the screen was light gray. Centered at the top of each receipt, written in black text in 18-point font, were the two attributes: the meal category and the day, with the meal category listed on the line directly above the day. The receipts were flashed on the screen, one at a time, for a total of 40 receipts, and the order of receipts was randomized for each participant. Each receipt appeared on the screen for 0.75 s followed by a blank screen for 0.50 s. We randomly selected a target day and target meal for each participant to appear with a fixed relative frequency of .25 (i.e., 10 flashes out of 40). We selected this frequency because it is approximately midway between the presumed ignorance prior of the day attribute (1/7, or .14) and the meal attribute (1/3, or .33). Thus, if participants anchor on ignorance priors suggested by the queried attributes, they should overestimate the true frequency of the target meal, and they should underestimate the true frequency of the target day. True frequencies of the other two possible meals and six possible days were randomly generated.

Following the learning phase, participants were asked to estimate the likelihoods of all 10 attributes (each of the three possible meals and seven possible days) in an order that was randomized for each participant. Participants in the probability condition were asked, "If a receipt is drawn at random from the set you just observed, what is your best estimate of the probability that the [meal / day] shown on the receipt is [breakfast / Sunday]? Please state your answer as a percentage (0%–100%)." In contrast, participants assigned to the proportion condition were asked, "In terms of the different types of meals in the set you just observed, what proportion of Joe's meals [were breakfasts] / [were eaten on Sunday]? Please state your answer as a percentage (0%–100%)."

Next, participants were asked about their relevant confidence and knowledge. The specific wording for the confidence item was "How confident are you in your recollection of the events you just provided estimates for?" The specific wording of the knowledge item was "How knowledgeable do you feel judging frequencies, proportions, and probabilities?" Both items were rated on a 7-point scale (1 = *not at all confident/knowledgeable*, 7 = *extremely confident/knowledgeable*).

Results

Response mode. We found no significant differences in likelihood judgments of the target attributes between participants in the proportion and probability conditions, nor did we find any differences in relative reliance on ignorance priors versus true frequencies ($p \sim .46$, two-tailed, across all analyses listed below). Thus, we pooled the data from these response modes.

Judgments of target events. The distributions of likelihood judgments for target meals and days are depicted in Figure 1. Although the true relative frequency of these events was .25, participants tended to underestimate the likelihood of target days and overestimate the likelihood of target meals, as we predicted, providing a median[4] judgment of .20 when queried about the target day and .30 when queried about the target meal. These median estimates, listed in Table 1, fall midway between true frequency and the corresponding ignorance priors and are statistically distinguishable from both. The median judgment of target day (.20) is significantly lower than the true frequency of .25 ($p = .005$, by Wilcoxon test) and significantly higher than the ignorance prior of .14 ($p < .001$). Likewise, the median judgment of target meal (.30) is significantly higher than the true frequency of .25 ($p = .003$, by Wilcoxon test) and significantly lower than the ignorance prior of .33 ($p < .001$).

Additional evidence of attribute-cued partitioning can be seen in an analysis of individual responses. The median within-participant discrepancy between judgments of the target meal and target day is .05, with a majority of participants (59%) reporting a higher likelihood for the target meal than the target day, 6% reporting equal likelihoods, and 35% reporting a higher likelihood for target day than target meal ($p < .001$, by sign test).

Judgments of relative frequency. To measure how accurately people learned overall relative frequency of the events, for each participant we computed the correlation between that individual's judged probabilities and the associated true frequencies of all 10 attributes (each of the seven possible days and three possible

[3] This cover story was inspired by Dougherty and Hunter (2003, Study 1).

[4] In most instances the mean results followed an analogous pattern to the median results. However, we chose to report medians and nonparametric statistics because the data were nonnormal (the Kolmogorov–Smirnov test yielded 2.8, $p < .001$, for target day and 1.6, $p < .01$, for target meal).

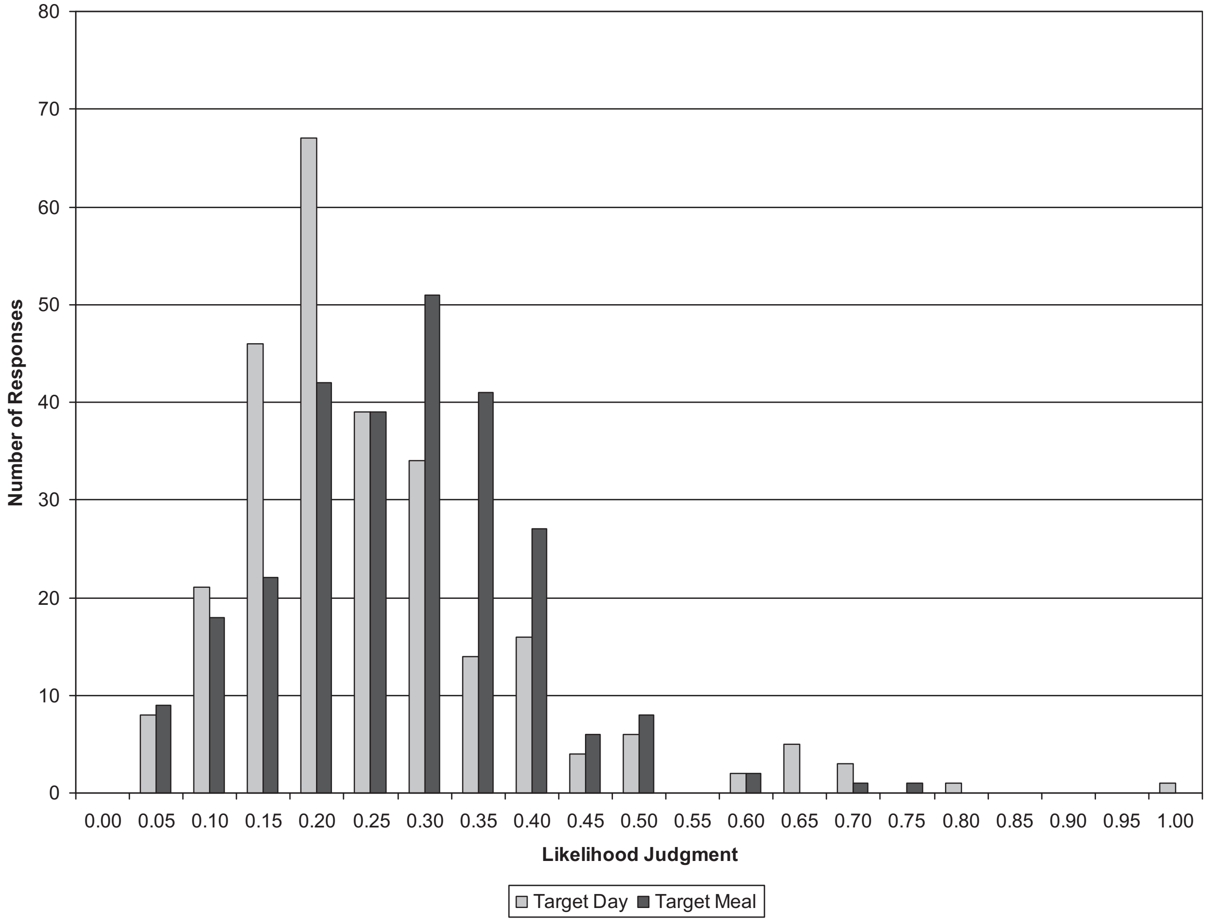

Figure 1. Distribution of likelihood judgments for target day and meal in Study 1 ($n = 267$).

meals). In order to obtain an estimate of relative frequency that was independent of the analyses of target judgments, we also computed the correlation between judged and true probabilities for only the 8 nontarget attributes. These measures reflect a general index of accuracy and learning of relative frequencies. The median correlation across participants was .75 for all 10 attributes and .81 for the 8 nontarget attributes, suggesting that participants were generally quite adept at learning relative frequencies and forming likelihood estimates.

Regression analysis. We next converted probabilities and relative frequencies to a log-odds metric and regressed likelihood judgments of all 10 attributes (seven days, three meals) on the ignorance prior and true frequencies, separately for each participant. The regression equation, derived in Appendix A from Fox and Rottenstreich's (2003) ignorance prior model, can be written as follows:

$$\ln R(F,A) = a + b_1 \ln\left[\frac{n_f}{n_a}\right] + b_2 \ln\left[\frac{f(F)}{f(A)}\right],$$

where $R(F, A)$ is the odds, derived from judged probabilities, of a focal event F (e.g., ate at restaurant on a Wednesday) rather than its alternative A (ate at restaurant on a day other than Wednesday); n_f and n_a are the number of qualitatively distinct levels that the focal and alternative events, respectively, can take on (e.g., 1 and 6 for days of the week); and $f(F)$ and $f(A)$ are the absolute frequencies in which the focal and alternative events, respectively, appeared during the learn-

Table 1
Results of Study 1

Sample	p(target day)	p(target meal)	p(meal) − p(day)	p(meal) > p(day) %	p(meal) = p(day) %	p(meal) < p(day) %
Overall ($N = 267$)	.20	.30	.05	59*	6	35
Low competence ($n = 145$)	.20	.30	.05	65*	3	32
High competence ($n = 122$)	.23	.26	.02	52.5	9	38.5

Note. Columns 2, 3, and 4 report median judged probabilities.
* $p < .001$ (one-tailed) by sign test.

Table 2
Median Unstandardized Regression Coefficients for All Studies

Study	Frequency	Ignorance prior	Intercept	R^2
Study 1				
Overall (n = 267)	0.46*	0.36*	−0.13*	0.61*
Low competence (n = 145)	0.42*	0.41*	−0.11	0.62*
High competence (n = 122)	0.53*	0.26*	−0.14*	0.61*
Study 2				
Moderate salience (n = 83)	0.72*	0.15	0.00	0.76*
Cognizant subset (n = 40)[a]	0.97*	0.01	0.00	0.78*
Low confidence (n = 21)	0.71*	0.15	0.00	0.78*
High confidence (n = 19)	1.06*	−0.05	0.00	0.75*
Enhanced salience (n = 84)	0.43*	0.50*	0.00	0.67*
Low confidence (n = 55)	0.46*	0.50*	0.00	0.66*
High confidence (n = 29)	0.40*	0.40*	0.00	0.71*
Study 3				
No-object-bar (n = 73)	0.80*	0.04	0.00	0.69*
Retrieval-only (n = 70)	0.50*	0.44*	0.00	0.71*
Encoding + retrieval (n = 86)	0.53*	0.48*	0.00	0.80*

[a] Participants who reported exactly four shapes and two colors.
* $p < .025$ (one-tailed) by sign test.

ing phase. Thus, for each participant, we obtained an intercept term and two separate regression coefficients: one coefficient that represented the influence of the ignorance prior and one that represented the influence of observed frequency.

Median results accord with our major predictions (see top section of Table 2). Participants relied on both observed frequency ($B = 0.46$, $p < .001$) and the attribute-cued ignorance prior ($B = 0.36$, $p < .001$). We note that the median intercept term of −.13 was small but statistically significant ($p < .001$). This value, although not predicted by the ignorance prior model (see Appendix A), reflects a small downward bias in all probabilities.[5] Overall, the fit of this model is good (median $R^2 = .61$).

Competence effects. We were interested in whether people exhibit more accurate judgments and less reliance on ignorance priors when they feel more confident in their memory of the events and more knowledgeable judging frequencies. The median confidence judgment was 3.0, and the median knowledge judgment was 4.0. These items formed a reliable scale (Cronbach's $\alpha = .74$), so we averaged the two items to form a composite measure of self-rated competence judging likelihood.

We first looked at the effects of competence on judgments of target events. Participants who rated themselves as less competent exhibited greater reliance on attribute-cued ignorance priors: The Spearman correlation between competence and the within-participant difference score for median ratings of the target meal and target day was −.12 ($p = .028$). To further probe this effect, we conducted a median split on the competence measure. As shown in Table 1, the median within-participant difference among those who rated themselves as highly competent was .02, whereas the median difference among those who rated themselves as less competent was .05, significantly larger ($p = .024$, by Kruskal–Wallis test).

We examined the regression results to see whether self-reported competence moderated the weights afforded ignorance priors versus true frequency. Because we were running separate regressions for each individual participant, it was not possible to statistically interact the competence measure with the independent variables within each regression. Rather, we again relied on a median split of the composite competence measure. Median coefficients for the low- and high-competence groups are listed in Table 2 and reveal significantly less reliance on true frequency ($p = .022$, by Kruskal–Wallis test) and greater reliance on ignorance priors ($p = .048$, by Kruskal–Wallis test) among those who rated themselves as less competent compared with those who felt more competent.

Discussion

Study 1 provides evidence that people rely on ignorance priors implied by the number of distinct values that the target attribute could potentially take on, even in a controlled learning environment in which frequency information is readily available. Regression analyses provide further evidence that participants placed appreciable weight on ignorance priors when forming likelihood judgments. In addition, participants relied more on the ignorance prior and less on their observation of true frequencies if they felt less competent assessing likelihood. Finally, Study 1 shows that none of these results are diminished among participants asked to judge relative frequency rather than probability.

The finding that participants place more weight on ignorance priors and less weight on observed frequencies when they feel less competent assessing likelihood could be viewed as rational Bayesian behavior in which participants adopt attribute-cued ignorance priors as prior probabilities and treat their impressions of relative likelihood as less diagnostic when they feel less competent. Likewise, the negative correlation between feelings of competence and reliance on ignorance priors could be accommodated by Fiedler and Armbruster's (1994) information-loss model, according to which judged frequencies of events regress toward an even distribution over all possible events the noisier or more imperfect the learning. For Study 1, both Bayesian and information-loss models predict that likelihood judgment will lie between true frequency and the ignorance prior, closer to true fre-

[5] To put this in perspective, the bias in judged probability implied by an intercept of −.13 varies from a maximum of .03 (in cases in which the ignorance prior model predicts a response of $p = .50$) to .01 (in cases in which the ignorance prior model predicts a response of $p = .10$ or .90).

quency as the judge's confidence in what has been learned increases (either because one's impressions are more diagnostic of true frequencies or because less information has been lost).

In the present article, we propose that the weight afforded the nature versus number of events is affected not only by participants' relative confidence in these sources of information but also by the relative salience of these cues, as in other information integration tasks (cf. Anderson, 1981). The question arises, however, how partition salience might interact with confidence. We surmise that reliance on ignorance priors reflects an automatic tendency to incorporate salient cues, whether or not these cues are perceived to be relevant. In this respect ignorance priors may contaminate likelihood assessment much like judgmental anchors contaminate magnitude estimation. Prior research has demonstrated that although the impact of self-generated anchors can be moderated by such factors as financial incentives, explicit injunctions for accuracy, or inducing participants to shake or nod their heads during judgment, the impact of more accessible externally provided anchors is not sensitive to such factors (Epley & Gilovich, 2001, 2005). In a similar vein, although we found that reliance on self-generated ignorance priors was sensitive to participants' perceived level of competence in Study 1, we predict that reliance on ignorance priors will be less sensitive to level of confidence if ignorance priors are made especially salient through an exogenous prompt. Indeed, just as highly accessible, externally provided anchors can contaminate judgments of even the most confident experts (Englich, Mussweiler, & Strack, 2006; Northcraft & Neale, 1987; Wright & Anderson, 1989), particularly salient ignorance priors may contaminate likelihood judgments of even the most confident or knowledgeable individuals. In the following study, we explore this potential interaction between perceived competence and salience of attribute-cued ignorance priors.

Study 2: Interaction of Confidence and Partition Salience

Studies 2 and 3 differ from Study 1 in two important respects. First, we move from a state space for which we expected participants to have a priori notions concerning the number of possible events (there are seven days in a week and three meals in a day) to a state space that is constructed in an ad hoc manner so that participants should have no prior expectations: Participants observed objects that could take on one of four different shapes in one of two different colors. Constructing an artificial state space allows us to experimentally control not only true frequencies and ignorance priors but also the salience of attribute-cued partitions. Second, we constructed a perfect correlation among the attributes (one color was uniquely associated with a single shape) so that we could ask participants to judge the probability of a particular target event by asking about different attributes of the same event. This allows for a stronger test of partition dependence, which, strictly speaking, refers to the tendency for judged likelihood of a particular event to vary with the relative accessibility of alternative partitions.

Participants viewed a distribution of four colored shapes that flashed one at a time on a computer screen with varying frequencies and then judged the probabilities of each shape and color. We manipulated the salience of attribute-cued partitions by either including an explicit visual cue that displayed all possible objects in a matrix by shape and color (cf. Fiedler & Armbruster, 1994, Experiment 1) or asking participants to report the different shapes and colors they had observed and then reminding them of their responses at time of judgment. We also elicited self-reported confidence and knowledge assessments as in the previous study.

The present account yields two testable predictions. First, we predict that more weight will be afforded to the ignorance prior relative to observed frequency (and partition dependence will therefore be more pronounced) when attribute-cued partitions are made more salient via an externally provided visual cue than when these partitions are made less salient via an internally generated list of attributes. Second, the impact of self-rated confidence on relative weights of these cues (and the impact of self-rated confidence on partition dependence) will diminish when partitions are more salient.

Method

Participants. Participants were 94 Duke University students and 73 University of North Carolina students, who were recruited to participate in a set of unrelated computer studies lasting a total of approximately 1 hr and for which they were paid $10.

Design and procedure. The experiment was administered on laptop computers in a private meeting room in the student union of each university. We used a between-participants design in which individuals were randomly assigned to one of two salience conditions: a moderate salience condition ($n = 83$) and an enhanced salience condition ($n = 84$).

Participants were told that they would view a distribution of multiattribute objects that would be quickly flashed on the center of the screen in a random order and instructed to try to get an intuitive sense of how often each object appeared. They were also told that after viewing the objects they would be asked to judge the probabilities of events they observed. The design features of this study are depicted in Appendix B. Each of the four objects in the set had two manipulated attributes: shape and color. Specifically, the objects were approximately .75 in. (1.9 cm) across and took on one of four shapes (triangle, circle, square, or diamond) and one of two colors (black or gray). The background was white. We expected that for most participants the shape attribute would suggest a fourfold partition and ignorance prior of 1/4, whereas the color attribute would suggest a twofold partition and ignorance prior of 1/2.

The relative frequencies of the four objects were determined at random, subject to the constraint that each of the four objects would appear a minimum of 5% of the time. Also, one target object was randomly selected for each participant to be the only one of its kind having a particular shape and color; this target object appeared with a fixed frequency of .35. We selected this frequency because it is approximately midway between the ignorance priors of the target shape (.25) and the target color (.50).

All participants viewed the set of objects that flashed 40 times, with each object appearing on the screen for .25 s followed by a blank screen for .10 s. The order in which objects appeared was randomized separately for each participant. Participants in the enhanced salience condition were presented with an object-bar at the top of the screen to remind them of the possible attribute levels, as depicted in Appendix B. We displayed the object-bar throughout the experiment (during instructions, learning, and judgment phases), which we expected to provide a very salient visual cue of the different attribute levels.[6] The position of the two colors and four shapes within the object-bar were randomized for each participant in this condition.

Participants in the moderate salience condition were never shown an object-bar or instructed on the possible attributes but were instead asked, after viewing the objects, how many shapes and colors they had seen: "Thinking about the set of objects you just observed and the characteristics of those objects, how many distinct [shapes / colors] did you observe?"

[6] We note that our use of this object-bar throughout the entire task in the enhanced salience condition resembles the method used by Fiedler and Armbruster (1994, Experiments 1 and 2) if one were to combine their encoding-split and response-split conditions.

Table 3
Results of Studies 2 and 3

Study	p(target shape)	p(target color)	p(color) − p(shape)	p(color) > p(shape) %	p(color) = p(shape) %	p(color) < p(shape) %
Study 2						
Moderate salience (n = 83)	.30	.30	.00	46	20	34
Cognizant subset (n = 40)[a]	.30	.30	.00	45	20	35
Low confidence (n = 21)	.30	.40	.05	57	14	29
High confidence (n = 19)	.30	.25	.00	32	26	42
Enhanced salience (n = 84)	.34	.45	.10	65*	11	24
Low confidence (n = 55)	.30	.45	.10	65*	13	22
High confidence (n = 29)	.35	.45	.10	65*	7	28
Study 3						
No-object-bar (n = 73)	.30	.30	.00	33	27	40
Retrieval-only (n = 70)	.35	.50	.10	63**	14	23
Encoding + retrieval (n = 86)	.30	.50	.10	63**	17	20

Note. Columns 2, 3, and 4 report median judged probabilities.
[a] Participants reporting seeing exactly four shapes and two colors.
* $p < .025$ (one-tailed) by sign test. ** $p < .001$ (one-tailed) by sign test.

Next, we asked all participants to judge the probabilities of the six attributes (each of the four possible shapes and two possible colors) in an order that was randomized for each participant. For example, we asked, "If an object were drawn at random from the exact same distribution you just observed, what is the probability that the object is [a square / gray]?" For participants in the moderate salience condition only, a sentence appeared at the bottom of the screen during probability elicitation that was designed to remind them of the attribute levels they had reported seeing: "You reported observing X shapes and Y colors." Finally, all participants were asked to report their perceived domain knowledge and confidence in their memory using the same items as in Study 1.

Results

Salience of attributes. When participants in the moderate salience condition were asked how many distinct shapes and colors they had seen, the median number of shapes reported was 4 (SD = 0.79) and the median number of colors was 2 (SD = 0.76). Of the 83 participants in this condition, 40 reported having seen exactly four shapes and two colors, and a significant majority reported having seen four shapes or two colors or both (72 of 83, $p < .001$). Thus, as expected, participants in the moderate salience condition appear to have encoded the attribute levels fairly accurately, though the correct values were obviously not as transparent as in the enhanced salience condition. In order to provide a comparable test of partition dependence across experimental conditions, we restricted our analyses in the moderate salience condition (except as indicated) to the subset of participants who reported seeing four shapes and two colors (i.e., the same partitions provided in the enhanced salience condition). However, results for this subset (n = 40) did not differ markedly from results for the rest of the moderate salience condition (n = 83), as can be seen by comparing the first two rows of Table 3.

Judgments of target events. The distributions of judged probabilities for target shape and color are provided in Figures 2A and 2B for the moderate salience condition (overall and subset, respectively) and Figure 2C for all participants in the enhanced salience condition. It is evident from these Figures that partition dependence is strong in the enhanced salience condition and greatly diminished in the moderate salience condition. Participants in the enhanced salience condition reported a median judged probability of .45 when they were queried about the target color and a median judged probability of .34 when they were queried about the target shape ($p < .001$, by Wilcoxon test). In striking contrast, participants in the moderate salience condition exhibited no significant partition dependence when queried about the two attributes of the same target event. The median participant reported a judged probability of .30 for the target shape and .30 for the target color (p = .49, by Wilcoxon test), with a median within-participant discrepancy of 0, significantly less than the corresponding .10 discrepancy in the enhanced salience condition (p = .009, by Kruskal–Wallis test).

As shown in the top panel of Table 3, median estimates for shape were quite close to the true frequency in both conditions, whereas median estimates for color were midway between the true frequency and ignorance prior for the enhanced salience condition only. In the enhanced salience condition, the median judgment of target color (.45) was significantly higher than the true frequency of .35 ($p < .001$, by Wilcoxon test) and significantly lower than the ignorance prior of 1/2 ($p < .001$), but the median judgment of target shape (.34) was not significantly less than the true frequency of .35 (p = .29). In the moderate salience condition, the median judgment of target color (.30) was not significantly different than the true frequency of .35 (p = .39, by Wilcoxon test), and the median judgment of target shape (.30) was marginally lower than the true frequency of .35 (p = .099).

A similar pattern holds in the analysis of individual participants (see the last three columns of Table 3). In the enhanced salience condition, a significant majority (65%) of participants judged the probability of the target color to be greater than that of the target shape ($p < .005$, by sign test), with only 11% judging the probabilities of the target color and shape to be equal. In the moderate salience condition, nearly twice as many participants (20%) judged the probabilities of the target attributes to be equal, with fewer than half (45%) of participants judging the probability of the target color to be greater than that of the target shape.

Judgments of relative frequency. We again computed the correlation for each participant between his or her probability judg-

A

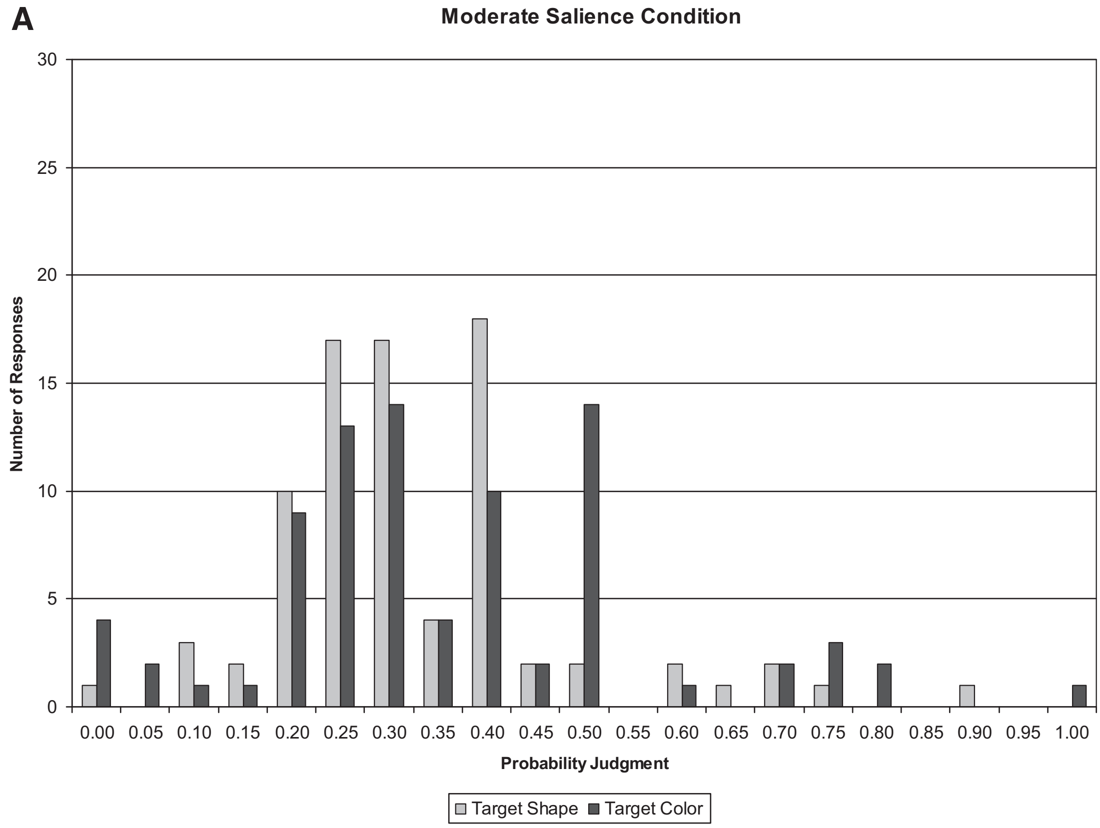

B

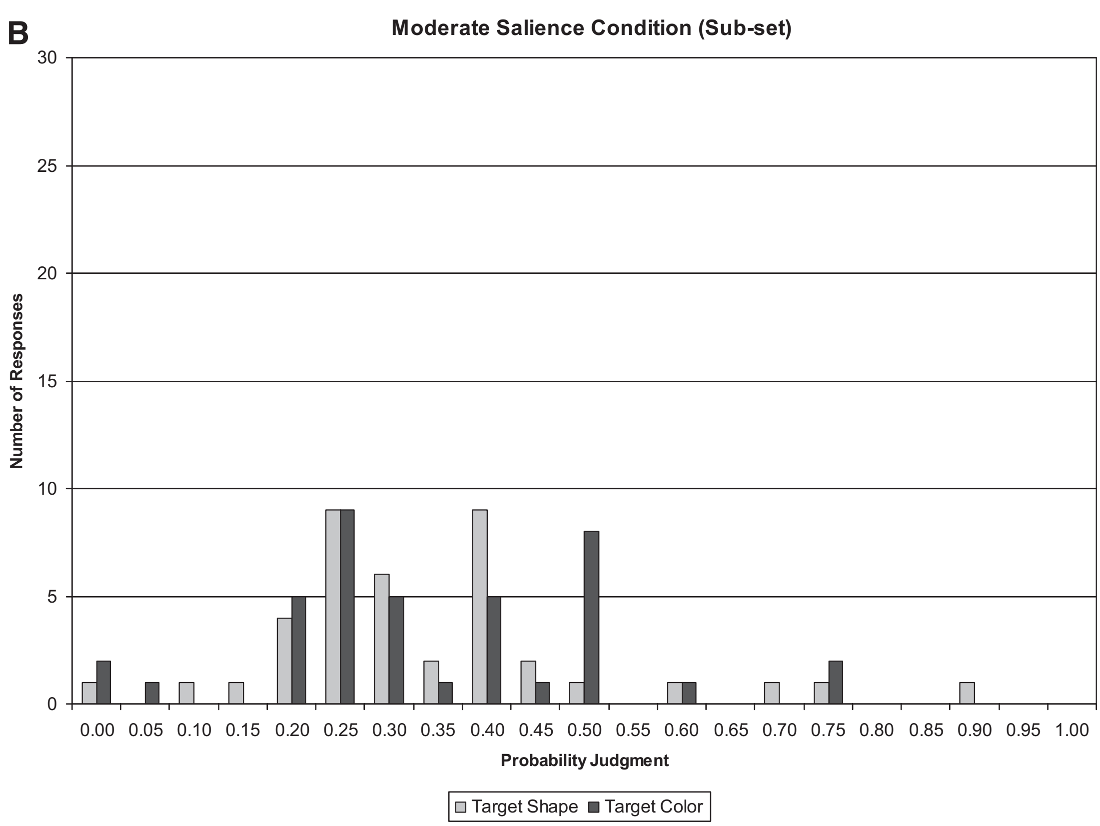

Figure 2. A: Distribution of probability judgments for target shape and color for the moderate salience condition of Study 2 (n = 83). B: Distribution of probability judgments for target shape and color for the moderate salience condition of Study 2 (n = 40 participants who reported seeing exactly four shapes and two colors). C: Distribution of probability judgments for target shape and color for the enhanced salience condition of Study 2 (n = 84).

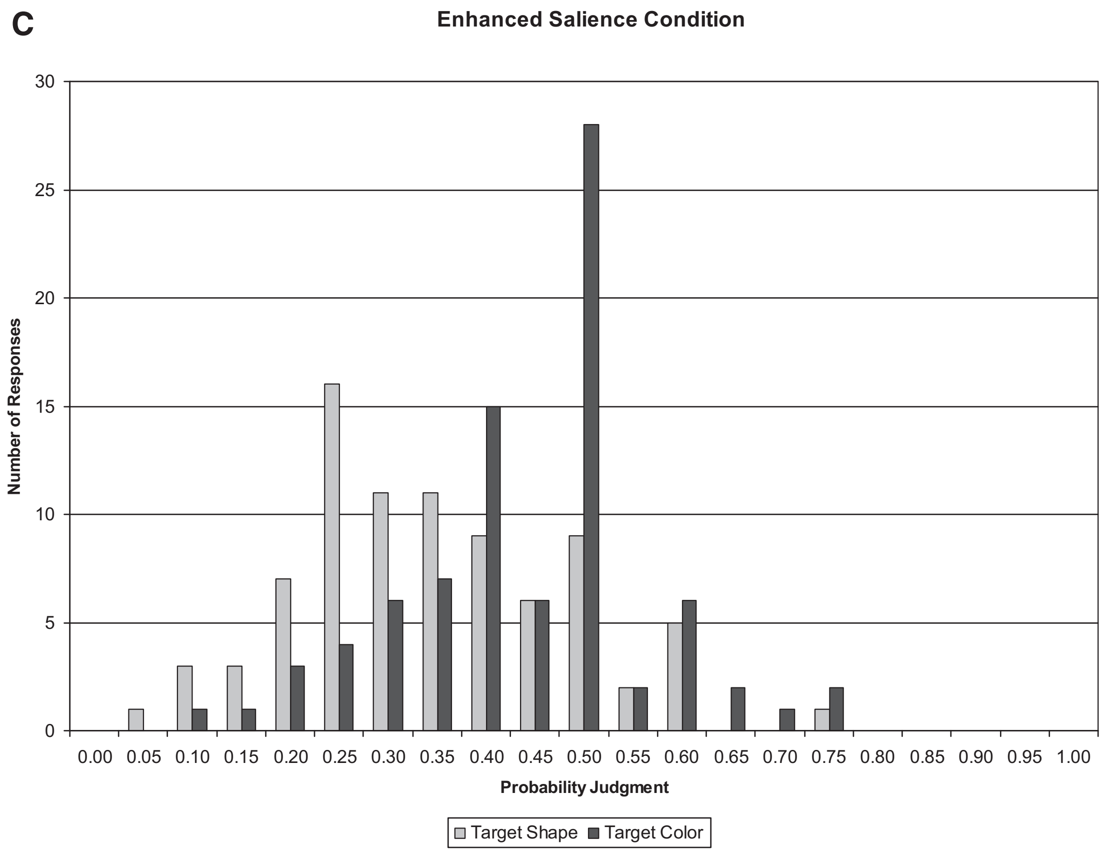

Figure 2 (*continued*).

ments and the relevant objective probabilities for all six attributes. We also computed a second measure that was independent of the analyses of target judgments by computing the correlation of judged and true probabilities for only the four nontarget attributes. We computed these correlations separately for each experimental condition. The median correlation for all six attributes was .85 for the moderate salience condition and .73 for the enhanced salience condition, a modest but significant difference ($p = .013$, two-tailed). This difference disappeared when the nontarget judgments were excluded: The median correlation was .96 for the moderate salience condition and .91 for the enhanced salience condition ($p = .20$, two-tailed). Thus, once again participants learned relative frequencies very well despite the intrusion of attribute-cued ignorance priors for target judgments in the enhanced salience condition.

Regression results. As in Study 1, we converted the data to log-odds and ran individual-level regressions using all six judgments made by each participant (four shapes, two colors). Results are listed in the middle section of Table 2. Participants in the enhanced salience condition appeared to rely on both the ignorance prior ($B = 0.50$) and observed frequency ($B = 0.43$), as both coefficients were significantly greater than zero ($p < .001$, by sign test) and were not different from each other ($p = .71$, by Wilcoxon test). Participants in the moderate salience condition, in contrast, relied significantly more on observed frequency ($B = 0.97$) than the ignorance prior ($p < .005$, by Wilcoxon test) and in fact did not put any appreciable weight on the ignorance prior ($B = 0.01$, *ns*). Note that the intercept term was not significant in either condition (median of 0.00). Thus, unlike in Study 1, we see no general upward or downward bias in assessed probability, only attribute-cued partition dependence in the enhanced salience condition. Importantly, reliance on ignorance prior versus true frequency was significantly different between conditions. The median weight afforded the ignorance prior was significantly lower in the moderate salience condition than in the enhanced salience condition ($p = .007$, by Kruskal–Wallis test). On the other hand, true frequency had more influence on judged probability in the moderate salience condition than in the enhanced salience condition ($p = .002$, by Kruskal–Wallis test). Overall, the regression results show an effect of both ignorance prior and observed frequency for the enhanced salience condition, but only an effect of frequency for the moderate salience condition.

Confidence effects. Our second prediction concerned the interaction of perceived competence with salience of attribute-cued ignorance priors. We expected that when attribute levels were less salient, reliance on ignorance priors would be more sensitive to people's confidence in their memory of the events and perceived knowledge judging frequencies. The median confidence judgment was 3.0, and the median knowledge judgment was 4.0. Unlike Study 1, these items did not form a reliable scale (Cronbach's $\alpha = .63$), so we were unable to combine them into an overall measure of perceived competence. Internal tests revealed that the knowledge variable had no effects on any results in the current sample, so we focused our analyses on self-reported confidence in memory.

There were no differences in reported levels of confidence in memory between the moderate salience condition (*Mdn* = 3.0, *SD* = 1.35) and the enhanced salience condition (*Mdn* = 3.0, *SD* = 1.18). This indicates that the manipulation of attribute salience did not perturb levels of confidence, and hence differences in partition dependence between the two conditions cannot be explained by confidence. Examining the effects of confidence on judgments of target events within each condition, we found that participants in the enhanced salience condition did not exhibit a relationship between confidence in memory and partition dependence: The Spearman correlation between confidence and the within-participant difference score for median ratings of the target color and target shape was −.06 (p = .29). In the moderate salience condition, however, confidence did have an effect: The Spearman correlation between confidence and the within-participant difference score for median ratings of the target color and target shape was −.29 (p = .036).

To further probe this interaction effect, we conducted a median split on the confidence measure to create low- and high-confidence groups for each condition. As shown in Table 3, the median within-participant difference between judged probabilities of target color and shape among those who rated themselves as highly confident in their memory in the moderate salience condition was .00, whereas the median difference among those who rated themselves as less confident was .05, significantly larger (p = .05, by Kruskal–Wallis test). In the enhanced salience condition, Table 3 shows that the median within-participant difference was .10 for both highly confident participants and those who rated themselves as less confident (p = .42, by Kruskal–Wallis test). Thus, as expected, confidence significantly affected the degree of partition dependence in target judgments for participants in the moderate salience condition only.

We also examined the regression results to see how self-reported confidence affected weights afforded ignorance priors versus true frequency, again using a median split of the confidence variable to compare unstandardized coefficients within each experimental condition. Median coefficients for the low- and high-confidence groups are listed in Table 2 and reveal significantly more reliance on true frequency (p = .04, by Kruskal–Wallis test) and less reliance on the ignorance prior (p < .06, by Kruskal–Wallis test) among those in the moderate salience condition who rated themselves as more confident versus less confident.[7] For the enhanced salience condition, the relatively equal weight on frequency and the ignorance prior did not vary significantly with level of confidence (p > .40 for both comparisons).

Discussion

Study 2 provides further evidence that judged probabilities are influenced by both observed frequencies and ignorance priors. We demonstrated attribute-cued partition dependence using a simplified ad hoc state space in which attribute levels were perfectly correlated. Partition dependence was strong and consistent when alternative partitions of the state space were made salient through the presence of an object-bar, but partition dependence was not found when alternative partitions were less salient. In addition, we observe that the influence of the ignorance prior in the enhanced salience condition was stronger for the color attribute than for the shape attribute; we return to this result in the General Discussion.

Perhaps most interesting, Study 2 demonstrates that partition dependence is more pronounced among those who are less confident in their memory for true frequencies only when alternative partitions are less salient. Likewise, regression analysis revealed that greater confidence is associated with more reliance on true frequency and less reliance on ignorance priors only when partitions are less salient. Confidence had no effect on the degree of partition dependence or on the weights afforded observed frequency versus ignorance priors in the enhanced salience condition. This pattern of results supports our view that reliance on ignorance priors often reflects the simple psychological tendency to incorporate cues that are salient, even though such cues may not always be relevant. Ignorance priors can thus resemble anchors that contaminate judgment of even highly confident or knowledgeable individuals.

The finding that partition salience influences the impact of ignorance priors prompts the question of whether salient cues (a) interfere with *encoding* of relative frequencies during the learning phase or (b) bias *retrieval* or expression of frequencies at the time of elicitation. Hastie and Park (1986) argued that the potential exists both for memory to bias on-line judgment (biased encoding) and for on-line judgment to bias memory (biased retrieval). The studies we have presented so far cannot distinguish among these possibilities; for example, in Study 2 the object-bar was presented during both learning and probability elicitation for the enhanced salience condition. However, Fiedler and Armbruster (1994) investigated encoding and retrieval processes in the context of a different judgment phenomenon, the category-split effect (a manifestation of explicit subadditivity; see Rottenstreich & Tversky, 1997). In one of their studies, they showed participants a distribution of five objects: four geometric symbols, one of which appeared in two different orientations. Some participants were verbally told of the four symbols (without mentioning that one symbol appeared in two different orientations) and asked to estimate the frequency of the four symbols; others were graphically presented with all five objects and asked to estimate the frequency of the five objects. The authors observed that the sum of estimates for different orientations of an object (e.g., "pentagons pointing up" and "pentagons pointing down") was greater than the single estimate of the object's overall frequency (e.g., "pentagons") and that graphically presenting the objects and their different orientations during likelihood elicitation (retrieval) was sufficient for this category-split effect to emerge.

In Study 3, we investigate whether the phenomenon of attribute-cued partition dependence will likewise emerge when alternative partitions are salient only during probability elicitation (retrieval), or whether the effect also requires alternative partitions to be salient during learning (encoding). If partition dependence requires partitions to be salient during both encoding and retrieval, it may suggest that partitions influence judgment by defining the categories into which data are organized as they are learned. If partition

[7] We also examined the median within-participant difference between weight on the ignorance prior and weight on observed frequency for participants in the moderate salience condition. The difference was marginally higher for the high-confidence group relative to the low-confidence group (p < .06).

dependence occurs even when partitions are salient only during retrieval, it would lend additional support for our notion that salient ignorance priors serve as prominent cues that can contaminate judgment.

Study 3: Encoding Versus Retrieval

In Study 3, we replicated the method of Study 2 with a first group of participants submitted to the enhanced salience condition in which an object-bar was displayed both during the learning and the probability elicitation phases. A second group of participants was presented with an object-bar only during the elicitation of probabilities. Finally, a third group of participants was not presented with an object-bar at all.

As with the previous study, we expect attribute-cued partitions to be much less salient when the object-bar does not appear so that partition dependence will be less pronounced in the condition with no object-bar than in the condition in which an object-bar is present throughout the task. Moreover, if partition dependence reflects reliance on salient ignorance priors that intrude much like anchors whenever they are accessible, we would expect to observe significant partition dependence even when the object-bar appears only at the time of judgment. If that occurs, we would expect to observe relatively accurate judgments in the condition in which no object-bar is present, consistent with automaticity of frequency encoding (Hasher & Zacks, 1984), but should see partition dependence in other conditions in which attributes are made salient at some point during the task.

Method

Participants. We recruited 229 Duke University undergraduates to participate in a series of unrelated experiments lasting approximately 20 min and for which they were paid \$5.

Design and procedure. The experiment was administered on laptop computers in a private meeting room in the student union. The design features of this study and stimuli are exactly the same as Study 2 except as indicated below.

Participants were assigned at random to one of three experimental conditions that varied only in terms of whether the visual object-bar used in the enhanced salience condition of Study 2 was presented throughout the task or only during the elicitation of probabilities. Participants assigned to the encoding-plus-retrieval condition (n = 86) saw the object-bar throughout the instructions, presentation of flashed objects, and elicitation of judged probabilities, exactly as in the enhanced salience condition of Study 2. Participants assigned to the retrieval-only condition (n = 70) saw the object-bar only at the time of judgment after they had observed the presentation of objects. The object-bar was introduced to participants in this condition at the probability elicitation phase with the statement: "Each object you saw could have been one of four shapes shown above (triangle, circle, square, diamond) and one of two colors (gray or black)." Finally, participants assigned to the no-object-bar condition (n = 73) were never presented with an object-bar and were given no instruction concerning the possible attribute levels at any point during the study.

Results

Judgments of target events. The distributions of judged probabilities for target shapes and colors are depicted in Figure 3A for the no-object-bar condition, Figure 3B for the retrieval-only condition, and Figure 3C for the encoding-plus-retrieval condition. Results are consistent with the notion that reliance on ignorance priors requires partitions to be salient only during retrieval. Participants in the encoding-plus-retrieval condition were strongly influenced by attribute-cued ignorance priors, replicating the results of the enhanced salience condition of Study 2. As shown in the bottom section of Table 3, the median participant in this group reported a judged probability of .30 when queried about the target shape and .50 when queried about the target color (p < .001, by Wilcoxon test). Similarly, participants in the retrieval-only condition reported a median judged probability of .35 for shape and .50 for color (p < .001, by Wilcoxon test). The median within-participant discrepancy between judgment of the target color and shape was .10 in both the retrieval-only and encoding-plus-retrieval conditions, which did not differ significantly (p = .49, two-tailed, by Kruskal–Wallis test). As with Study 2, partition dependence seems to be stronger for the color attribute in this study.

In contrast, participants in the no-object-bar condition exhibited no significant partition dependence when queried about the two attributes of the same target event: The median participant reported a judged probability of .30 for the target shape and .30 for the target color (p = .16, by Wilcoxon test), with a median within-participant discrepancy of .00, significantly less than the corresponding discrepancy in the retrieval-only condition (p < .001, by Kruskal–Wallis test) and encoding-plus-retrieval condition (p < .001, by Kruskal–Wallis test). However, the median target shape judgment of .30 across all three conditions, which is midway between the true frequency of .35 and the ignorance prior of .25, suggests that there was some partition dependence for the shape attribute; we return to this result in the General Discussion.

A similar pattern holds at the level of individual participants (see Table 3). In both the retrieval-only and the encoding-plus-retrieval conditions, nearly two thirds of participants judged the target color to be more likely than the target shape, whereas fewer than one quarter of participants judged the target color to be less likely than the target shape (p < .001, by sign test for both conditions). However, in the no-object-bar condition, about as many participants judged the target color to be more likely than the target shape as the converse (p = .32, by sign test). In sum, judgments of the target shape and color suggest that participants who were shown the object-bar only at the time of elicitation responded similarly to those who were shown the object-bar throughout the experiment, but these groups differed substantially from those who were never shown the object-bar.

Judgments of relative frequency. The median correlation between judged and true probability for all six attributes was .81 for the no-object-bar condition, .75 for the retrieval-only condition, and .79 for the encoding-plus-retrieval condition. For the nontarget attributes, the median correlation between judged and true probability was .97 for the no-object-bar condition, .93 for the retrieval-only condition, and .95 for the encoding-plus-retrieval condition. As in the previous studies, participants showed impressive learning of relative frequencies, and median correlations were not statistically distinguishable between conditions ($p \sim .26$ across all comparisons). Echoing previous research on probability biases in learning paradigms, which has documented similarly high aggregate correlations between judged and true probability (e.g., Windschitl et al., 2002), we found that the biased patterns of attribute-

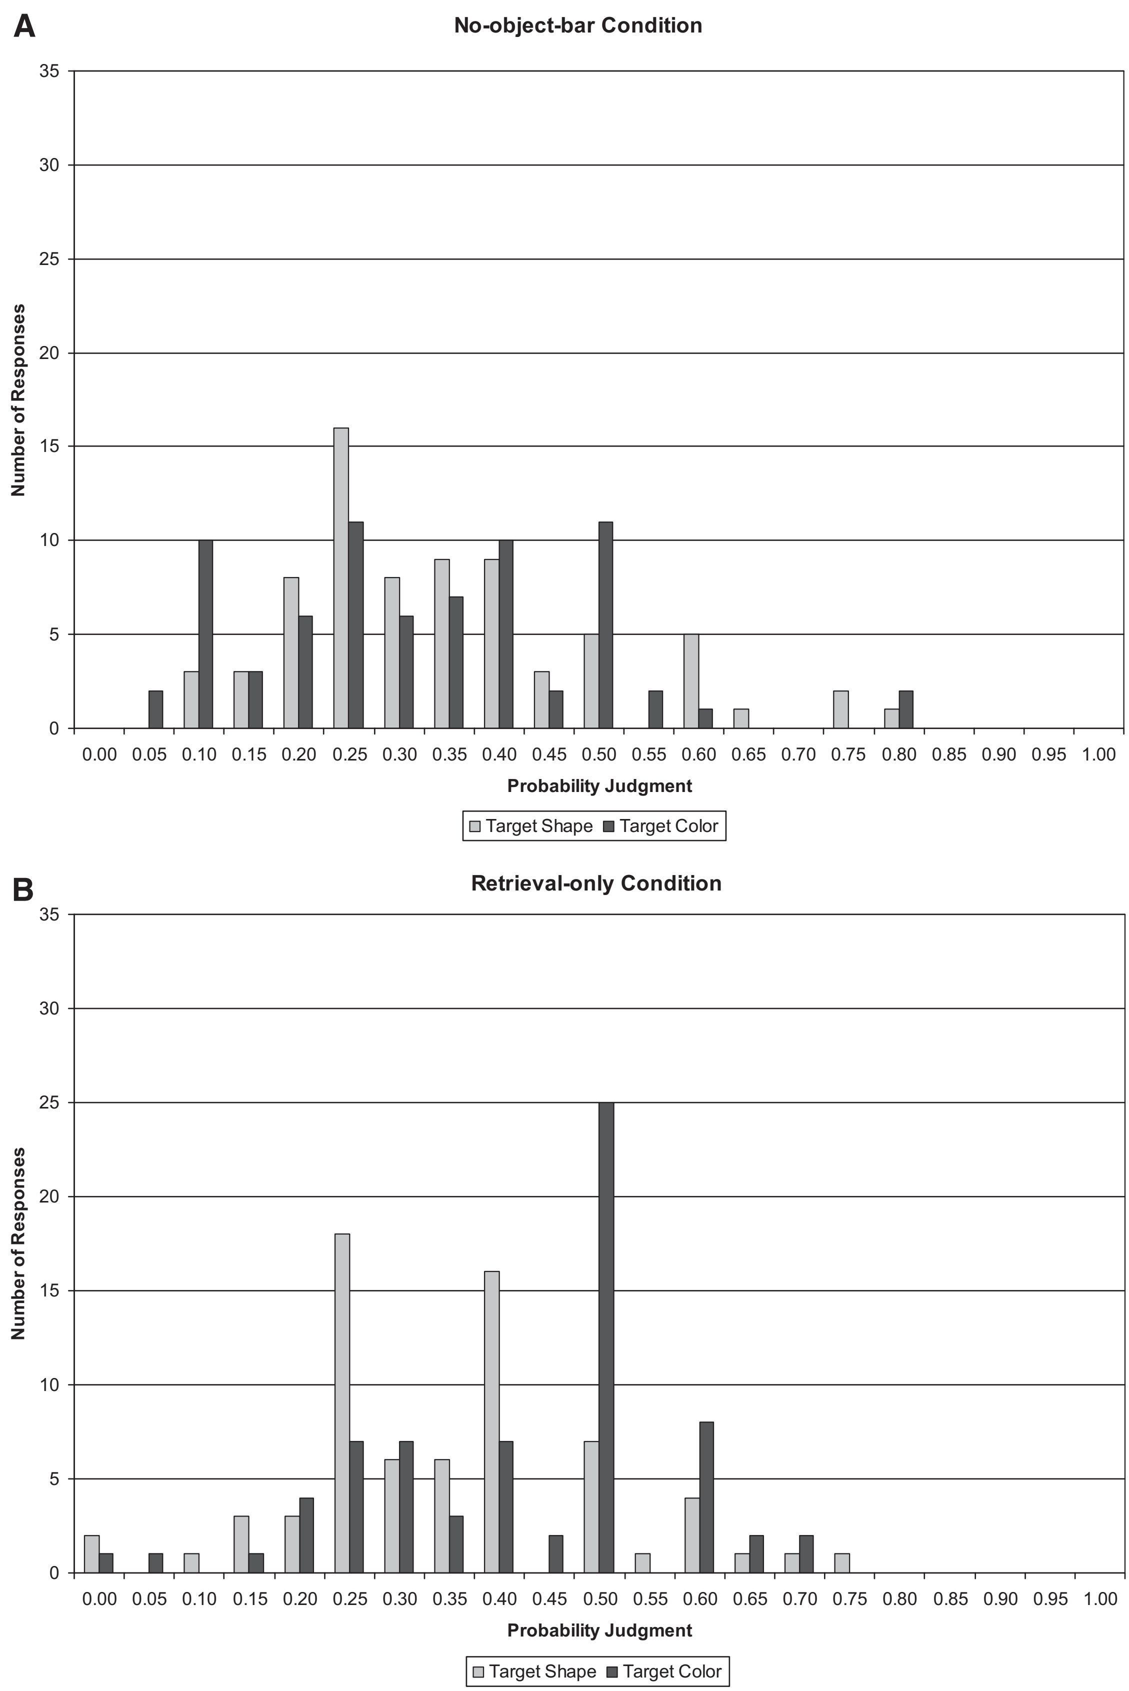

Figure 3. A: Distribution of probability judgments for target shape and color for the no-object-bar condition of Study 3 (n = 73). B: Distribution of probability judgments for target shape and color for the retrieval-only condition of Study 3 (n = 70). C: Distribution of probability judgments for target shape and color for the encoding-plus-retrieval condition of Study 3 (n = 86).

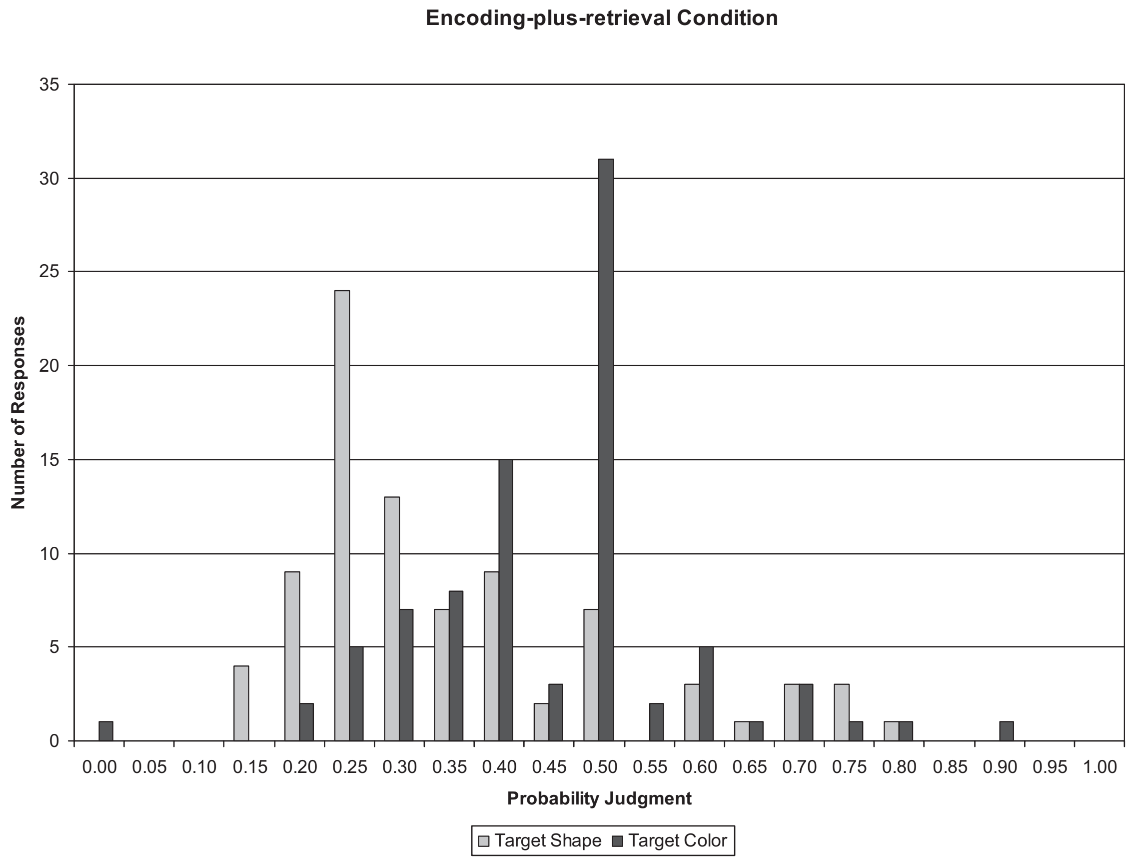

Figure 3 (*continued*).

cued partition dependence occurred in the context of otherwise quite accurate judgments of relative frequency.

Regression results. We again converted the data to log-odds and regressed, for each participant, judged probability on attribute-cued ignorance priors and true frequency. The results of this analysis are displayed in Table 2. For the no-object-bar condition, the median coefficient of the ignorance prior was negligible (B = 0.04, p = .30, by sign test), whereas the median coefficient for frequency was substantial (B = 0.80, p < .001). Within the retrieval-only and encoding-plus-retrieval conditions, however, participants gave roughly equal weight to the ignorance prior and true frequency; the two coefficients in both conditions were significantly different from zero (p < .001). Moreover, the coefficients were statistically indistinguishable from each other between the retrieval-only and encoding-plus-retrieval conditions (p = .87 (two-tailed) for frequency and p = .63 (two-tailed) for ignorance prior by Kruskal–Wallis test). However, the median results in the no-object-bar condition differed significantly from the other two groups. Participants in the no-object-bar condition relied significantly less on the ignorance prior (p < .001, by Kruskal–Wallis test) and significantly more on true frequency (p < .01, by Kruskal–Wallis test) relative to the other two conditions. Finally, we note that as in Study 2, the median intercept was precisely 0 for all three conditions (see Table 2). This suggests that unlike Study 1, there was no overall bias in judged probabilities other than the bias toward attribute-cued ignorance priors.

Discussion

Study 3 provides new evidence that intrusion of the ignorance prior occurs at the time of retrieval (or probability construction) rather than at the time of learning. In the no-object-bar condition, the only systematic inaccuracy we uncovered was a slight tendency for shape judgments to be biased toward the shape ignorance prior, which may have been more accessible without an object-bar than the color attribute. Evidently, participants in this condition relied primarily on retrieval of automatically encoded frequency information (Hasher & Zacks, 1984). However, in the encoding-plus-retrieval and retrieval-only conditions, judged probabilities were systematically biased toward the corresponding ignorance prior. This pattern of results is consistent with our interpretation that relatively accurate encoding of frequency information can subsequently be contaminated by salient cues concerning the number of possible events. It is interesting to note that, in learning studies of phenomena as diverse as partition dependence (the present article) and the category-split effect (Fiedler & Armbruster, 1994), it appears that judgmental bias requires biasing information to be present only during retrieval.

General Discussion

Substantial past research has suggested that people are adept at automatically registering counts of events that they have observed (e.g., Hasher & Zacks, 1984), whereas at the same time, relative frequency and proportion estimations may be biased toward 1/2 or other salient reference points (e.g., Hollands & Dyre, 2000). In the present article, we attempted to integrate these disparate literatures by proposing that people typically rely on both a somewhat accurate recollection of true frequencies and an ignorance prior derived from the number of elementary events into which the judge subjectively partitions the state space. This can give rise to partition dependence, in which judged probabilities vary systematically with the number of elementary events into which the state space is subjectively parsed.

Previous work has demonstrated that alternative partitions can be made salient either by explicitly providing alternative sets of exclusive and exhaustive events for participants to assess (Fox & Clemen, 2005) or by varying the language of the probability query (Fox & Levav, 2004; Fox & Rottenstreich, 2003; Smithson & Segale, 2006). In this article, we observe for the first time that alternative partitions can also be cued by drawing attention to different attributes that vary in the number of levels they can take on. We present a simple hypothetical example (the auto dealership cited in the introduction) and three studies that provide the first evidence of (a) within-participant partition dependence, (b) that is cued by attributes, and (c) which occurs in a controlled learning environment. Study 1 used a familiar state space (based on days of the week and meals of the day) for which attribute-cued partitions would be naturally accessible to participants and showed that reliance on ignorance priors is affected by perceived competence judging frequencies. Study 2 used an unfamiliar ad hoc state space (objects that varied in shape and color) in which attribute-cued partitions were made more salient to some participants via an externally provided visual cue and showed that salience of the partition moderates the impact of confidence on partition dependence. In particular, when partitions were made especially salient through presentation of an object-bar, partition dependence was not related to feelings of confidence. Finally, Study 3 revealed that partition dependence does not require that attribute-cued ignorance priors be salient during learning (encoding) but rather that it occurs even when they are made available only at the time of probability elicitation (retrieval).

Attributes as Categories Versus Properties

It is notable that reliance on ignorance priors in Studies 2 and 3 seemed to be stronger for the color attribute than for the shape attribute. In fact, median judgments for target shape were quite accurate. This observation accords with the results of Barsalou and Ross (1986), who found that participants were highly accurate in reporting the number of instances presented from superordinate categories (e.g., birds) but that their estimates of the frequencies of properties (e.g., red) were less accurate (see also Conrad, Brown, & Dashen, 2003). Johnson, Peterson, Yap, and Rose (1989) further found that an item's membership in a particular category must be salient for participants to accurately judge the relative frequency of that category. Similarly, Freund and Hasher (1989) found that participants were more accurate estimating a category's frequency if they had been informed in advance that they would be asked to judge the frequency of that category (see also Burton & Blair, 1991; Conrad, Brown, & Cashman, 1998).

We suspect that participants in our studies had an easier time accessing the number of distinct shapes (a category that is processed more spontaneously) than the number of distinct colors (a property that may be less salient without external prompting), and hence, they demonstrated a stronger impression of the relative frequency of shapes than of color. We surmise that shape acted like a superordinate category, more readily identifying objects than color, a property that refines the identification of the object.

To What Extent Is Reliance on Ignorance Priors Bayesian?

Our main purpose in this article has been to investigate the impact of the nature of events (true frequency) versus the number of events (attribute-cued ignorance priors) on likelihood judgment. Though our theoretical framework is descriptive in nature, one could argue that some instances of reliance on ignorance priors might be interpreted as normatively defensible Bayesian behavior. For example, the results of Study 1 showing that the impact of ignorance priors decreased when people were more confident in what they had learned might be consistent with a Bayesian account whereby people begin with prior beliefs that receipts are evenly distributed over all three meals and all seven days and perceive their impressions of relative likelihood to be more diagnostic of true frequencies as confidence increases.

The results of Study 2 showing that reliance on the ignorance prior is not sensitive to confidence when partitions are especially salient is less readily accommodated within a Bayesian framework without additional assumptions about the diagnostic value of salience. In addition, manifestations of partition dependence outside of a learning paradigm may be difficult to rationalize. We note from previous research that reliance on ignorance priors can influence judgment even in situations in which they imply peculiar Bayesian priors. For instance, Fox and Rottenstreich (2003) found that people asked to judge the probability that "next week, the highest temperature of the week will occur on Sunday" were biased toward 1/7 (a reasonable prior implied by a sevenfold partition by days of the week); however, people asked to judge the probability that "the temperature on Sunday will be higher than every other day next week" were biased toward 1/2 (a peculiar prior implied by a twofold partition by whether or not the statement is true). Likewise, the car dealership example in the introduction shows that even though participants can observe that "cars sold by Carlos" and "coupes" refer to the same event, they nevertheless seem to adopt different ignorance priors to these queries rather than a single compound prior, a particularly striking pattern of description dependence when replicated within-participant (see Footnote 2).

Even if one views reliance on ignorance priors as Bayesian, the critical questions remain what factors influence the priors that people happen to invoke and what factors influence their weight in judgment. In this respect the cognitive account advanced here contributes to the literature on judgment under uncertainty by demonstrating that (a) prior beliefs are often influenced by the number of qualitatively distinct levels that a target attribute may take on, (b) the relative weight afforded these priors versus learned information is a function of factors including confidence in learning and salience of attributes, and (c) salience can moderate the impact confidence so that when an ignorance prior is especially salient it can influence judgment regardless of one's level of confidence.

The tendency to rely on salient ignorance priors irrespective of confidence is interesting given the very high median correlations between judged and true frequencies across all three learning studies in which information was provided to participants rather quickly. This supports the notion that people often automatically encode relative frequency information quite accurately (Hasher & Zacks, 1984), and thus reliance on ignorance priors may serve to bias otherwise accurate judgment toward arbitrary reference points. In this sense, ignorance priors can contaminate responses much like judgmental anchors. Some prescriptive suggestions for improving accuracy of judged probabilities that may be influenced by ignorance priors have been offered by Fox and Clemen (2005). In addition, Clemen and Ulu (2006) have offered a procedure to counteract bias toward ignorance priors in the context of a Bayesian model of probability judgment.

Future Research

What partitions do people naturally invoke? Previous demonstrations of partition dependence have manipulated the relative accessibility of alternative partitions using explicit prompts. Fox and Clemen (2005) asked participants to consider state spaces that were explicitly divided into different sets of events to be judged, as is common practice in decision analysis. Fox and Rottenstreich (2003) used linguistic formulations that either highlighted a target event against its complement or highlighted a set of interchangeable events (e.g., days of the week). Likewise, Fox and Levav (2004) used alternative formulations of common probability puzzles to prompt different subjective partitions of state spaces that gave rise to a different proportion of correct responses.

In the present studies, we directed participants' attention to alternative attributes of various state spaces, but we did not otherwise explicitly prompt consideration of alternative partitions. In the restaurant study, participants seem to have spontaneously invoked a sevenfold partition for days of the week and a threefold partition for meals of the day. In the colored shapes studies, we simultaneously made both shape-based and color-based partitions accessible for some participants through presentation of an object-bar. Although the manipulation of partition salience in the present studies was arguably more subtle than were the manipulations used in previous articles, participants in our studies seemed to naturally invoke the partition cued by the attribute on which they were focusing their attention, and they were therefore biased toward the corresponding ignorance prior.

We surmise from the present results and those of previous studies cited above that at least partial reliance on ignorance priors is nearly ubiquitous. The default is to rely on a twofold "case" partition {target event occurs; target event does not occur} and an ignorance prior of 1/2, because these are always accessible when a target event is identified. However, in various contexts (such as the restaurant study) attributes may naturally cue "class" partitions by the number of levels of an attribute or the accessibility of a class of *n* interchangeable events. For ad hoc state spaces in which there are no such natural classes of interchangeable events, we suspect that the partitions that people invoke are influenced by how they construe the topology of the state space. Indeed, even in the no-object-bar condition of Study 3 in which we observed no partition dependence, we suspect that many participants relied on ignorance priors that were the same for the target color and shape. We note that judged probabilities in that condition (Mdn = .30 for both target shape and color) tended to lie below the true value of .35, biased in the direction of an integrated object-cued ignorance prior of 1/4 (recall that four unique objects were presented to participants). More generally, we surmise that for state spaces containing a small number of events, integrated object-cued partitions may take precedence over attribute-cued partitions, especially when attributes can be integrated into whole objects or distinctive events (cf. Brase, Cosmides, & Tooby, 1998).[8] Further investigation of what kinds of partitions people spontaneously invoke in natural settings is an important avenue for future research.

In addition, it would be interesting to examine naturally occurring instantiations of attribute-cued partition dependence. In many social contexts, natural attributes such as ethnicity or gender may suggest a number of discrete categories and therefore different ignorance priors. For instance, if a large set of candidates for a job includes several women and men who are African American, Asian American, Hispanic, and White, the judged probability that a woman is hired may be biased toward 1/2, whereas the judged probability that an Asian American is hired may be biased toward 1/4.

The role of confidence in memory. One of the more provocative findings of the present investigation is how readily people shift away from their relatively accurate impressions of frequency toward an ignorance prior when one is accessible. When attribute-cued partitions are only moderately salient, a critical factor influencing reliance on observed frequencies is an individual's feeling of confidence or knowledge. Indeed for the self-generated partitions in Study 1 and moderate salience condition of Study 2, partition dependence was stronger among those who rated themselves as less confident in memory. These findings are interesting in that they might suggest that subjective perceptions of one's memory are sufficient to affect probability judgment in learning environments, which builds on earlier work showing that individual differences in objective working memory can have important effects on probability judgment (Dougherty, Gettys, & Ogden, 1999; Dougherty & Hunter, 2003; Sprenger & Dougherty, 2006). This said, participants exhibited very accurate judged relative frequencies across the studies, but the correlation between judged relative frequency and reported confidence in memory was only .10 ($p \sim .05$) in Study 1 and .15 ($p < .05$) in Study 2. Thus, it appears that people are largely unaware of the extent to which they have accurately encoded frequencies and may be prone to discounting their automatic counts. Future research might explore whether accuracy can be improved (at least in situations in which attribute-cued partitions are not especially salient) by boosting people's confidence in their ability to automatically encode frequency information.

Conclusion

Throughout this article we have contrasted judgment based on the number of possible events (reliance on ignorance priors) with judgment based on the nature of those events (direct observation of relative frequency). Analysis of the balance struck between these two processes suggests that people both encode frequencies relatively accurately and are influenced by ignorance priors and that

[8] We note that although the state space in the Fletcher auto sales example on the first and second pages of the introduction contained only four relevant events, the attributes of model and salesperson could not be readily integrated into a single whole object or distinctive event.

use of ignorance priors does not vary as a function of confidence when partitions are very salient. In addition, the present article provides the first demonstration of partition dependence within-participant, using attribute-cued partitions, and in a controlled learning environment. We hope that the present investigation will encourage further attempts to integrate the distinct yet complimentary literatures on frequency encoding, categorization, and judgment under uncertainty.

References

Alba, J. W., Chromiak, W., Hasher, L., & Attig, M. S. (1980). Automatic encoding of category size information. *Journal of Experimental Psychology: Human Learning and Memory, 6,* 370–378.

Anderson, N. H. (1981). *Foundations of information integration theory.* New York: Academic Press.

Attneave, F. (1953). Psychological probability as a function of experienced frequency. *Journal of Experimental Psychology, 46,* 81–86.

Barsalou, L. W., & Ross, B. H. (1986). The roles of automatic and strategic processing in sensitivity to superordinate and property frequency. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 116–134.

Brase, G. L., Cosmides, L., & Tooby, J. (1998). Individuation, counting, and statistical inference: The role of frequency and whole-object representations in judgment under uncertainty. *Journal of Experimental Psychology: General, 127,* 3–21.

Burton, S., & Blair, E. (1991). Task conditions, response formulation processes, and response accuracy for behavioral frequency questions in surveys. *Public Opinion Quarterly, 55,* 50–79.

Clemen, R. T., & Ulu, C. (2006). *Interior additivity and subjective probability assessment of continuous variables.* Manuscript submitted for publication.

Conrad, F. G., Brown, N. R., & Cashman, E. R. (1998). Strategies for estimating behavioural frequency in survey interviews. *Memory, 6,* 339–366.

Conrad, F. G., Brown, N. R., & Dashen, M. (2003). Estimating the frequency of events from unnatural categories. *Memory & Cognition, 31,* 552–562.

Dougherty, M. R. P., Gettys, C. F., & Ogden, E. E. (1999). MINERVA-DM: A memory processes model for judgments of likelihood. *Psychological Review, 106,* 180–209.

Dougherty, M. R. P., & Hunter, J. (2003). Probability judgment and subadditivity: The role of working memory capacity and constraining retrieval. *Memory & Cognition, 31,* 968–982.

Englich, B., Mussweiler, T., & Strack, F. (2006). Playing dice with criminal sentences: The influence of irrelevant anchors on experts' judicial decision making. *Personality and Social Psychology Bulletin, 32,* 188–200.

Epley, N., & Gilovich, T. (2001). Putting adjustment back in the anchoring and adjustment heuristic: Differential processing of self-generated and experimenter-provided anchors. *Psychological Science, 12,* 391–396.

Epley, N., & Gilovich, T. (2005). When effortful thinking influences judgmental anchoring: Differential effects of forewarning and incentives on self-generated and externally provided anchors. *Journal of Behavioral Decision Making, 18,* 199–212.

Erlick, D. (1964). Absolute judgments of discrete quantities randomly distributed over time. *Journal of Experimental Psychology, 67,* 475–482.

Fiedler, K., & Armbruster, T. (1994). Two halfs may be more than one whole: Category-split effects on frequency illusions. *Journal of Personality and Social Psychology, 66,* 633–645.

Fox, C. R. (1999). Strength of evidence, judged probability, and choice under uncertainty. *Cognitive Psychology, 38,* 167–189.

Fox, C. R., & Clemen, R. T. (2005). Subjective probability assessment in decision analysis: Partition dependence and bias toward the ignorance prior. *Management Science, 51,* 1417–1432.

Fox, C. R., & Levav, J. (2004). Partition-edit-count: Naive extensional reasoning in judgment of conditional probability. *Journal of Experimental Psychology: General, 133,* 626–642.

Fox, C. R., & Rottenstreich, Y. (2003). Partition priming in judgment under uncertainty. *Psychological Science, 14,* 195–200.

Fox, C. R., & Tversky, A. (1998). A belief-based account of decision under uncertainty. *Management Science, 44,* 879–895.

Freund, S. J., & Hasher, L. (1989). Judgments of category size: Now you have them, now you don't. *American Journal of Psychology, 102,* 333–353.

Gigerenzer, G. (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, England: Wiley.

Gigerenzer, G. (1996). On narrow norms and vague heuristics: A reply to Kahneman and Tversky (1996). *Psychological Review, 103,* 592–596.

Gilovich, T., Griffin, D., & Kahneman, D. (Eds.). (2002). *Heuristics and biases: The psychology of intuitive judgment.* New York: Cambridge University Press.

Hasher, L., & Zacks, R. T. (1984). Automatic processing of fundamental information: The case of frequency occurrence. *American Psychologist, 39,* 1372–1388.

Hastie, R., & Park, B. (1986). The relationship between memory and judgment depends on whether the judgment task is memory-based or on-line. *Psychological Review, 93,* 258–268.

Hollands, J. G., & Dyre, B. P. (2000). Bias in proportion judgment: The cyclical power model. *Psychological Review, 107,* 500–524.

Johnson, M. K., Peterson, M. A., Yap, E. C., & Rose, P. M. (1989). Frequency judgments: The problem of defining a perceptual event. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 126–136.

Johnson-Laird, P. N., Legrenzi, P., Girotto, V., Legrenzi, M. S., & Caverni, J. P. (1999). Naive probability: A mental model theory of extensional reasoning. *Psychological Review, 106,* 62–88.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases.* Cambridge, England: Cambridge University Press.

Kahneman, D., & Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review, 10,* 582–591.

Koehler, D. J. (1996). A strength model of probability judgments for tournaments. *Organizational Behavior and Human Decision Processes, 66,* 16–21.

Koehler, D. J. (2000). Probability judgment in three-category classification learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26,* 28–52.

Naveh-Benjamin, M., & Jonides, J. (1986). On the automaticity of frequency coding: Effects of competing task load, encoding strategy, and intention. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 378–386.

Northcraft, G. B., & Neale, M. A. (1987). Experts, amateurs, and real estate: An anchoring-and-adjustment perspective on property pricing decisions. *Organizational Behavior and Human Decision Processes, 39,* 84–97.

Parducci, A., & Wedell, D. H. (1986). The category effect with rating scales: Number of categories, number of stimuli, and method of presentation. *Journal of Experimental Psychology: Human Perception and Performance, 12,* 496–516.

Rottenstreich, Y., & Tversky, A. (1997). Unpacking, repacking, and anchoring: Advances in support theory. *Psychological Review, 104,* 406–415.

Sheridan, T., & Ferrell, W. (1974). *Man-machine systems.* Cambridge, MA: MIT Press.

Smithson, M., & Segale, C. (2006). *Partition priming in judgments of imprecise probabilities.* Unpublished manuscript, Australian National University, Canberra, New South Wales.

Sprenger, A., & Dougherty, M. R. (2006). Differences between probability and frequency judgments: The role of individual differences in working memory capacity. *Organizational Behavior and Human Decision Processes, 99,* 202–211.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review, 90,* 293–314.

Tversky, A., & Koehler, D. J. (1994). Support theory: A nonextentional representation of subjective probability. *Psychological Review, 101,* 547–567.

Varey, C., Mellers, B., & Birnbaum, M. (1990). Judgments of proportion. *Journal of Experimental Psychology: Human Perception and Performance, 16,* 613–625.

Wickens, C. (1992). *Engineering psychology and human performance* (2nd ed.). New York: Harper-Collins.

Windschitl, P. D., Young, M. E., & Jenson, M. E. (2002). Likelihood judgment based on previously observed outcomes: The alternative outcomes effect in a learning paradigm. *Memory & Cognition, 30,* 469–477.

Wright, W. F., & Anderson, U. (1989). Effects of situation familiarity and financial incentives on use of the anchoring and adjustment heuristic for probability assessment. *Organizational Behavior and Human Decision Processes, 44,* 68–82.

Appendix A

The Ignorance Prior Model

To provide a more rigorous formulation of partition dependence, we describe Fox and Rottenstreich's (2003) formal model of partition dependence that generalizes support theory (Rottenstreich & Tversky, 1997; Tversky & Koehler, 1994) to incorporate reliance on ignorance priors.

In support theory, judged probability, $P(.)$, depends on the "support" or strength of evidence for a focal hypothesis, F (e.g., this person is a Democrat) relative to the support or strength of evidence for an alternative hypothesis, A (e.g., this person is not a Democrat). Denoting the support function by s, the judged probability that F will occur rather than A, is given by

$$P(F,A) = \frac{s(F)}{s(F) + s(A)}.$$

Letting $R(F, A) = P(F, A)/[1 - P(F, A)]$ be the odds for the focal hypothesis relative to the alternative hypothesis, we can rewrite the expression above as the following:

$$R(F,A) = \frac{s(F)}{s(A)}.$$

We can distinguish between support arising from the ignorance prior and support generated by evaluative methods by rewriting the expression above as

$$R(F,A) = \left[\frac{n_F}{n_A}\right]^{1-\lambda}\left[\frac{s^*(F)}{s^*(A)}\right]^{\lambda}.$$

Here n_f and n_a are the number of elements in the subjective partition corresponding to the focal and alternative hypotheses, respectively. For instance, for the Fletcher example in the introduction of this article, we assume that $n_f = 1$ and $n_a = 1$ for most people when the target attribute is models (coupe, sedan) and that $n_f = 1$ and $n_a = 3$ for most people when the target attribute is salesperson (Carlos, Jennifer, Damon, Sebastian). The values $s^*(F)$ and $s^*(A)$ quantify support generated through evaluative methods. For the present purposes, we assume that these values reflect participants' recall of the relative frequencies of each object. Finally, $0 \leq \lambda \leq 1$ is the weight given to evaluative assessments relative to the ignorance prior. As λ approaches 1, judged probability is based entirely on evaluative assessments; in contrast, as λ approaches 0, judged probability collapses to the ignorance prior. Note that whenever $\lambda > 0$ (positive weight is afforded the ignorance prior) judged probabilities will exhibit some degree of partition dependence.

Taking the log of both sides of the previous equation yields:

$$\ln R(F,A) = \beta_1 \ln\left[\frac{n_F}{n_A}\right] + \beta_2 \ln\left[\frac{s^*(F)}{s^*(A)}\right],$$

where $\beta_1 = 1 - \lambda$ is the weight afforded the ignorance prior and $\beta_2 = \lambda$ is the weight afforded evaluative assessments. If we assume that hypothetical support for a particular event is a power function of its raw frequency, f, we get

$$s^*(.) = [f(.)]^k,$$

where k is a scaling constant. For instance, $k < 1$ implies that support for an attribute is a concave function of the number of presentations of that attribute (for fuller discussions of scaling support, see Fox, 1999; Koehler, 1996; Tversky & Koehler, 1994). Thus, by regressing participants' judgments on the ignorance prior and on the objective frequencies of each object (all transformed to log-odds format), we can obtain a measure of the relative weight placed on the ignorance prior and on evaluative methods. If we also include an intercept term, we can detect whether there is residual bias in probability assessment:

$$\ln R(F,A) = \alpha + \beta_1 \ln\left[\frac{n_f}{n_a}\right] + \beta_2 \ln\left[\frac{f(F)}{f(A)}\right].$$

(*Appendixes continue*)

Appendix B

Experimental Design Features of Studies 2 and 3

Participants watched a set of multiattribute objects flash on the computer screen (one at a time) with varying frequencies and were then asked to judge the probability of these events. Each object could take on one of four shapes and one of two colors (see Figure B1). A target object was randomly selected for each participant to have a fixed frequency of .35 and take on a unique shape and color. For example, for a particular trial the target object might be a "black triangle," thus the other objects in the set would be a gray circle, gray square, and gray diamond, as shown below.

Figure B1.

The ignorance prior probability is $1/n$, where n is the number of distinct features of an event or object. Thus, the ignorance prior of a particular shape = 1/4, because n = 4 shapes; and the ignorance prior of a particular color = 1/2, because n = 2 colors (see Figure B2).

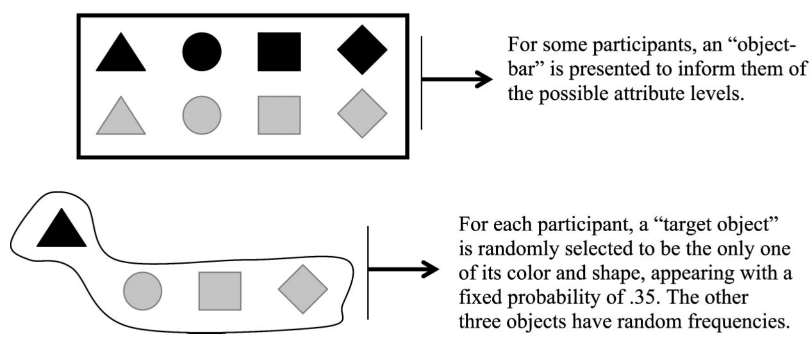

Figure B2.

Received December 30, 2005
Revision received June 9, 2006
Accepted June 15, 2006 ■