

MEMORY AND REPRESENTATIVENESS

Pedro Bordalo Katherine Coffman Nicola Gennaioli
Frederik Schwerter Andrei Shleifer*

December 11, 2019

Abstract

We explore the idea that judgment by representativeness reflects the workings of episodic memory. In our model, subjects assess the probability of a hypothesis given data in terms of the ease of retrieval of instances of that hypothesis when cued with the data. Retrieval is driven by a measure of similarity which exhibits contextual interference: a data-cue is less likely to retrieve a hypothesis that occurs in other data. As a result, retrieval and probability assessments are context dependent. In a new laboratory experiment on cued recall, participants are shown two groups of images with different distributions of colors. In line with the model's predictions, we find that i) decreasing the frequency of a given color in one group significantly increases the recalled frequency of that color in the other group, ii) cueing different dimensions for the same set of images entails different interference patterns and different probabilistic assessments. A calibration of the model yields a good quantitative fit with the data, highlighting the central role of contextual interference.

*Bordalo: University of Oxford; pedro.bordalo@sbs.ox.ac.uk. Coffman: Harvard Business School; kcoffman@hbs.edu. Gennaioli: Bocconi University and IGIER; nicola.gennaioli@unibocconi.it. Schwerter: University of Cologne; frederik.schwerter@uni-koeln.de. Shleifer: Harvard University; shleifer@fas.harvard.edu. We thank Rahul Bhui, Ben Enke, Paul Fontanier, Sam Gershman, Thomas Graeber, Spencer Kwon, Joshua Schwartzstein, seminar participants at Harvard University and MIT Sloan, and four anonymous referees for helpful comments. Gennaioli thanks the European Research Council (GA 647782) for financial support. Schwerter and Shleifer thank the Sloan Foundation for financial support.

1 Introduction

A vast literature in psychology shows that individuals' probabilistic judgments are vulnerable to systematic errors (see Benjamin 2018 for a review). An important body of evidence, ranging from base rate neglect to the disjunction fallacy, focuses on the representativeness heuristic, the tendency to judge as likely events that are merely representative. Representativeness captures “the degree to which [an event] is similar in essential characteristics to its parent population” (Kahneman and Tversky, KT 1972). In a well-known example, subjects are given a short description of an introverted man and are asked to rank, in order of likelihood, several different occupations (Tversky and Kahneman, TK 1974). Subjects tend to state that the introverted man is more likely to be a librarian than a salesman, even though there are vastly more salesmen than librarians. Being an “introvert”, the reasoning goes, makes the man more similar to a librarian than to a salesman, which causes an inflated assessment of likelihood.

But why do subjects use similarity in making probability judgments? Why should similarity sometimes cause them to think about unlikely events, so that judgments are incorrect? Why among introverted men do we think about the unlikely librarian instead of the more likely salesman? Why do we exaggerate the likelihood that an Irish person has red hair even though only 10% of them do (Bordalo et al., BCGS 2016)?

We address these questions starting from the idea that probability judgments are formed by selectively retrieving information from memory, cued by the problem at hand. In our model, a decision maker assesses the likelihood of different hypotheses in light of data d . In line with memory research (Kahana 2012), recall of a hypothesis h is driven by its similarity with the data d , and is subject to interference, so that hypotheses more similar to the data prevent less similar ones from coming to mind. When told someone is Irish – that is, when cued with the data $d = \text{Irish}$ – the decision maker selectively recalls the hair color h , in this case red, that is most similar to this data, and overestimates its likelihood.

Our main theoretical contribution is to formalize similarity between the data-cue and the hypothesis assessed, which constitute distinct and overlapping sets. Our formulation of similarity is inspired by Tversky's (1977) similarity metric, and captures the idea that two sets are more similar when the statistical correlation among their attributes is higher. In probabilistic judgments, this implies that a hypothesis h is more easily recalled when it is more likely to occur with the data d , but also when it is less likely to occur with other data $-d$. Similarity-based recall is then context dependent, in that recall about data d depends on the comparison data $-d$, sometimes causing neglect of likely events. For instance, when thinking about $d = \text{Irish}$, we find it harder to recall the modal hair color $h = \text{dark}$ because it occurs very often in other nationalities $-d$, which interferes with

its recall in the Irish group. We refer to this mechanism as contextual interference, since interference with hypothesis h is determined by the context provided by the alternative data $-d$.¹

Our second contribution is to present a novel experimental framework and a set of experimental results that provide evidence for this mechanism. The model yields testable predictions for how biases in probability judgments depend on the statistical associations between the data and hypotheses. To test these predictions, we run three (main) experiments with the following basic structure. Participants observe a sequence of 25 numbers and 25 decoy images (either shapes or words), which come in different colors. Among numbers, 15 are blue and 10 are orange. In one treatment, decoy images are all gray shapes so there is no color overlap with numbers. In another treatment, decoys are all blue words, which is also the modal color among numbers. Participants are then asked to: i) recall the numbers they saw, and ii) estimate the likelihood of different colors among numbers. We think of data as $d = \text{numbers}$, and the possible hypotheses, $h \in H$, as the colors.

In Study 1 we test contextual interference: we predict that the hypothesis $h = \text{orange}$ is deemed more likely when the decoy data $-d$ are blue words than when they are gray shapes, because blue words interfere with the recall of blue numbers. For the same reason, fewer blue numbers should be recalled in this case. In Study 2 we vary the strength of interference by adding orange words to the decoy data $-d$. As blue words are replaced with orange ones, retrieval of blue numbers is less interfered with and retrieval of orange numbers is more interfered with. Our model predicts that the assessed likelihood of hypothesis $h = \text{orange}$ should decrease with the number of orange words.

In Study 3 we test a key implication of our model of memory-based probabilistic judgments: descriptions of the judgment task that shape perceptions of similarity affect memory retrieval and thus probability judgments, even if they are normatively irrelevant. Specifically, we keep the color distribution of numbers and words fixed, but allow the images to differ along font size: either small or large. Crucially, all blue numbers are small, while all other items (orange numbers and blue words) are large. This ensures that, given $d = \text{numbers}$, the hypothesis $h = \text{orange}$ is equivalent to the hypothesis $h = \text{large}$. However, our model predicts that participants should assess the size distribution

¹The literature on similarity usually focuses on judgments of similarity between two objects of the same type, say similarity between faces (Tversky 1977, Kahana 2012). In contrast, we argue that to understand probabilistic judgments we must consider similarity between two sets: the data and the hypothesis. An example may help illustrate the difference. Consider the data “things that fly”. A standard approach would consider the similarity among two specific airplanes depending on their differences along certain characteristics. This is the approach followed also in probabilistic judgment models such as PROBEX. In our approach, instead, we ask how similar is the broad set “airplanes” to the broad set “things that fly”, which also includes birds and clouds, and thus how likely is “airplane” to be retrieved when cued with “things that fly”.

of numbers differently from its color distribution, even though the two coincide. The reason is that when thinking about size, contextual interference by the decoy data penalizes large numbers (which are orange), but when thinking about color it penalizes blue numbers (which are small).

We find strong and robust evidence consistent with these predictions. To assess how our model quantitatively accounts for the observed variation in beliefs about the frequency of orange numbers. We use the data from Study 1 and from a subset of treatments in Study 2 to calibrate the model. We then use the calibrated model to predict, out of sample, the other results in Studies 2 and 3, as well as in additional experiments we ran. This exercise confirms that contextual interference is a quantitatively important determinant of beliefs and that our model offers a good quantitative fit.

In our approach, the representativeness heuristic derives from a basic cognitive mechanism: context-dependent, similarity-based recall. As such, it differs from existing models of faulty probability judgments. One strand of the literature invokes limited memory, including sampling models (Sanborn and Chater 2016), the Minerva MD model (Dougherty et al. 1999) or the exemplar PROBEX model (Nilsson, Olsson, and Juslin 2005). The first two models are based on noisy recall, which yields some biases but not the overestimation of unlikely events. In the PROBEX model, the exemplar hypothesis of data d does not depend on different data $-d$, which makes this model unsuitable to account for how probability judgments depend on the decoy group.²

Bhatia (2017) shows that social stereotypes, such as introverted male librarians and red-haired Irish, can be spread through the media and personal interactions, through selective repetition of the more representative cases. This semantic association and repetition of unlikely traits becomes a source of accessibility in memory that may cause their overestimation. While this mechanism certainly contributes to biases in social contexts, it does not explain the source of these biases (why does media focus on these specific traits), and why biases are widespread even in abstract judgments such as those in our experiments.

²Other models of memory are motivated not by biases in probabilistic assessments but by evidence on recall, specifically by the fan effect (Anderson 1974) and a large body of evidence on interference going back to the early 20th century (Jenkins and Dallenbach 1924, Whitely 1927, McGeoch 1932, Underwood 1957, Keppel 1968). Anderson (1974) shows that concepts associated with more items are more difficult to remember in response to any specific cue, as evidenced by slower response times and a larger error rate in a recognition task. Similarly, in word-pair recall tasks, associating the same cue word with different target words in different lists reduces recall of both the pair learned first and the pair learned later (see Kahana 2012 for a review). Two broad approaches to this evidence have been proposed: associative models such as the Adaptive Control of Thought - Rational (ACT-R, Anderson and Reder 1999), and inhibition models such as inhibitory control in retrieval (Anderson and Spellman 1995). As in most existing models within the similarity framework, interference here occurs across types associated with a cue, which in our model corresponds to interference across different hypotheses h associated with data d . Instead, our model and experiments highlights the importance of contextual interference, namely the association of hypothesis h with some other data $-d$.

Several papers explicitly engage with the representativeness heuristic. Tenenbaum and Griffiths (2001) formalize representativeness based on Tversky and Kahneman (1983), but do not link it to similarity-based recall and interference. Our analysis provides a foundation for the representativeness heuristic in terms of interference in recall. More broadly, we show that interference in recall generates beliefs that amplify true differences between distributions, a property introduced in Gennaioli and Shleifer (GS 2010) and BCGS (2016) that we refer to as the kernel of truth.

Our recent work shows that the kernel of truth helps account for a wide range of experimental and field evidence on beliefs in several domains: GS (2010) illustrate its explanatory power for several puzzles in probabilistic judgments, such as the conjunction fallacy and base rate neglect, that were originally attributed to representativeness (KT 1972). BCGS (2016) show that the same principle helps explain measured beliefs about political groups as well as social stereotypes, while BCGS (2019) show that its predictive power extends to beliefs about own ability, tainted by group stereotypes. The kernel of truth also helps account both qualitatively and quantitatively for measured expectations of market participants in macroeconomics and finance (Bordalo, Gennaioli, Shleifer 2018, Bordalo, Gennaioli, Ma, Shleifer 2018, Bordalo, Gennaioli, La Porta, Shleifer 2019).

The results in this paper link the kernel of truth to memory mechanisms, deepen our understanding of these biases, and shed light on how beliefs are shaped by contextual cues. The model does not explain several other judgment biases that have been documented by KT (1972) and others, such as the law of small numbers or biases due to anchoring and to availability heuristics. We return to these in the conclusion.³

The paper is organized as follows. Section 2 describes the model and the experimental framework. Section 3 describes our first set of experiments in which representativeness is controlled by varying the comparison group. Section 4 presents the experiment in which the database is the same for all participants but recall is cued along different dimensions. Section 5 discusses our findings and Section 6 concludes.

³Growing literatures in both Neuroscience and Economics explore the role of memory for valuation and decision making, particularly the role of cues in selectively retrieving information. Bornstein et al (2017) find that choices among risky alternatives can be altered by showing subjects images that co-occurred with payoffs in the past, even if these images are totally uninformative. The authors argue that the subjects' behavior is consistent with a cued recall mechanism: whether the subjects respond to the cue is commensurate with the reinstatement in the brain's visual areas of patterns associated with the cue. Enke, Schwerter and Zimmerman (2019) give experimental subjects news about a hypothetical asset, associating certain images with good news and others with bad news. When assessing the asset's value, subjects overreact to good news associated with images that co-occurred with past good news. This result is in line with the model of memory-based choice of Bordalo, Gennaioli, and Shleifer (2019) in which context (an image) cues recall of past experiences associated with it. Several papers document that individuals' decisions are shaped by recall of past experiences: assessments of expected inflation are based on goods bought frequently (Cavallo, Cruces, Perez-Truglia 2017), and willingness to participate in the stock market depends on own past returns (Malmendier and Nagel 2011, Wachter and Kahana 2019).

2 Model and Experimental Framework

2.1 The Model

Our model captures the idea that probability judgments, such as assessing the probability distribution of hair color among the Irish, are formed by selectively retrieving relevant instances from memory. Retrieval of different hypotheses (hair colors) is cued by the data (the Irish), and is driven by their similarity. We show that Tversky’s (1977) model of similarity among sets yields a form of interference in recall that is critical to generating representativeness effects and several other predictions explored below.⁴

The memory database is described by a probability space with event space Ω and probability measure P , which summarizes a person’s experiences including their frequencies. Two random variables are defined on this space: H , which we think of as the hypotheses the DM seeks to assess, and D , which we think of as data. The task of the decision maker is to assess the probability of hypothesis $h \in H$ given data $d \in D$. In our running example, Ω is the universe of people, $D = \text{nationalities}$ and $H = \{\text{dark}, \text{light}, \text{red}\}$ are different hair colors.

A Bayesian retrieves all possible hypotheses h given d and computes the true conditional probability $P(h|d)$ using the measure P .⁵ Relative to this benchmark, in our model the data d cues retrieval of more similar hypotheses from memory, facilitating recall of certain hypotheses relative to others, and causing their probability in d to be inflated.

To formalize the similarity between data d and a hypothesis h we rely on Tversky’s (1977) similarity metric, which applies to objects characterized as collections of features. Given two objects a and b with sets of features A and B , respectively, the similarity of a to b is defined as:

$$S(a, b) = \alpha f(A \cap B) - \beta f(A/B) - \gamma f(B/A), \quad (1)$$

where $A \cap B$ is the subset of features that are common to both objects, A/B and B/A are the features that are present in one object but not in the other, $f(\cdot)$ is an increasing function, and α, β, γ are non-negative coefficients. According to Equation (1), the similarity between a and b is ceteris paribus higher when they have more features in common, and when each set has fewer distinct features not shared with the other.

⁴Tversky’s (1977) model of similarity has also been used to study the problem of categorization, which is distinct from what we consider here. Kemp, Bernstein, and Tenenbaum (2005) and Gershman (2017) show that testing whether two objects have the same data generating process reduces, under some assumptions, to a measure of non-Euclidean similarity similar to Tversky’s model.

⁵Formally, the decision maker computes the conditional probability by retrieving all elementary events ω consistent with each hypothesis h and data d , and by using the probability measure P to compute:
$$P(H = h|D = d) = \frac{\int_{\omega \in \Omega: H(\omega)=h, D(\omega)=d} dP(\omega)}{\int_{\omega \in \Omega: D(\omega)=d} dP(\omega)}.$$

Equation (1) can be applied to the context of probabilistic assessments by noting that data d and each possible hypothesis h are also characterized by features, namely the different possible values of the random variables under consideration. For instance, $d = \text{Irish}$ is the set of people who are Irish but can have different hair color (red, light, dark), while $h = \text{red}$ is the set of people who have red hair and may have different nationalities (Irish, Spanish, ...).

In this interpretation of Equation (1), the similarity between the set of red haired people and the set of Irish people depends on the size of three subsets. Similarity $S(d, h)$ between d and h increases in the size $P(d, h)$ of the set of people who are both Irish and have red hair. But similarity between d and h decreases in the size $P(d, -h)$ of the set of people who are Irish but do not have red hair, as well as in the size $P(-d, h)$ of the set of people who have red hair but are not Irish. Formally, the similarity between hypothesis h and data d is given by:

$$S(d, h) = \alpha f(P(d, h)) - \beta f(P(d, -h)) - \gamma f(P(-d, h)). \quad (2)$$

Similarity between hypothesis h and data d increases in their correlation: the more h and d occur together. In the previous example, for a given size of the red-haired Irish population $P(d, h)$, similarity between Irish and red hair is maximized when all Irish are red haired, $P(d, -h) = 0$, and when red hair occurs only among the Irish, $P(-d, h) = 0$.

Similarity is linked to retrieval through Luce's rule. That is, the probability of recalling h when cued with d is given by:

$$\tilde{P}(h|d) = \frac{e^{S(d, h)}}{\sum_{h' \in H} e^{S(d, h')}}. \quad (3)$$

When cued with d , a hypothesis h more similar to d is more easily retrieved. Retrieval is normalized, so that recall probabilities add to one. Normalization captures a form of interference whereby recall of more similar items inhibits that of less similar ones. The decision maker estimates the probability of h in d by measuring the frequency with which h comes to mind among all retrieved d members. Through the law of large numbers, Equation (3) also yields the probability that subjects on average attach to hypothesis h given data d .

According to (3), recall of h , and thus its assessed probability, follow the well known law of frequency. The more people are both Irish and have red hair (higher $P(d, h)$), the more likely is $d = \text{Irish}$ to cue retrieval of $h = \text{red hair}$. This effect is common to most memory-based models of probability assessments (Anderson and Reder 1999, Dougherty, Gettys, Ogden 1999, Kahana 2012).

The law of frequency is however subject to two kinds of interference. The first is

interference across different hypotheses. A person who has seen many dark haired Irish (high $P(d, -h)$) is less likely to recall red hair. This is due to both the β term in the similarity function (2) and the normalization of recall in Luce’s rule (3). Many memory models, including Minerva DM and sampling, display this form of interference, which tends to produce underestimation of unlikely traits.

The second, and critical, form of interference occurs across different data. For a decision maker who has seen few non-Irish people with red hair (which implies a low $P(-d, h)$), Irish is more associated with red hair, and retrieval of red haired Irish is increased. Conversely, when thinking about dark hair among $d = \text{Irish}$, Spaniards or Italians may come to mind because this hair color is much more common among these other nationalities $-d$. This interferes with recalling dark haired Irish. This form of interference flows back from hypothesis h to different data $-d$.

This second mechanism only operates when $\gamma > 0$, and implies that *distinctive* traits are overestimated, even if they are unlikely. We call this the “kernel of truth” property. It leads subjects to overestimate a hypothesis h that is more associated with the current data d relative to other data $-d$, even though the latter association is normatively irrelevant. Our experiment tests this prediction precisely by varying the normatively irrelevant “decoy” distribution $P(-d, h)$ while holding the relevant probabilities $P(d, h)$ and $P(d, -h)$ constant.⁶

Consider a convenient formulation in which $f(x) = \ln(c + x)$, where c is a non-negative constant, and $\beta = 0$.⁷ The assessed probability $\tilde{P}(h|d)$ then becomes:

$$\tilde{P}(h|d) \propto \frac{[c + P(h|d)P(d)]^\alpha}{[c + P(h|-d)P(-d)]^\gamma}. \quad (4)$$

Parameter α captures the strength of frequency-based recall. When α is high, likely hypotheses are oversampled from memory and their probability is inflated. For low α , the reverse holds: sampling depends less on $P(h|d)$, so that low likelihood hypotheses are oversampled across the board.

Parameter γ instead captures contextual interference, or the kernel of truth. If γ is high, when cued by data d we oversample hypotheses that are unlikely to occur with other data, i.e. with low $P(h|-d)$. The decision maker then inflates the probability of such distinctive hypotheses. The parameter c pins down a baseline probability of recall-

⁶Another difference from existing memory models is that in our setup encoding is perfect, not noisy as in the Minerva-DM of Dougherty, Gettys, Ogden (1999). Noisy encoding could be added to our model, but we stress that interference across groups relies on fairly accurate memories, in the sense that the co-occurrence of types and groups is correctly recorded. This is critical to obtain belief distortions that depend on the true features of the data.

⁷In our two-hypothesis setting, both β and α modulate the extent to which similarity $S(d, h)$ increases in the likelihood $P(d, h)$, so β is redundant.

ing a hypothesis, regardless of its frequency, which caps distortions when a hypothesis does not occur at all given the decoy data $-d$.

Equation (4) flexibly nests different kinds of judgments, both rational and biased. When $\alpha = 1$ and $c = \gamma = 0$, the model collapses to Bayesian inference. When instead $\alpha > 1$, and $c = \gamma = 0$ the model generates a form of over-reaction to data, in the sense of overestimation of hypotheses that are more likely given d , reminiscent of overconfidence (Moore and Healy 2008). Conversely, when $\alpha < 1$ the model generates a form of under-reaction to data. In the extreme case of $\alpha = 0$, all hypotheses are recalled with the same probability, so that unlikely ones are mechanically oversampled.

When $\gamma = 0$, the model fails to yield context-dependent beliefs and overestimation of distinctive hypotheses. This occurs only when $\gamma > 0$, so that assessments about d depend on the decoy distribution $-d$. This mechanism generates representativeness effects. In fact, if $\alpha = 1 + \gamma$ and $c = 0$, Equation (4) becomes:

$$\tilde{P}(h|d) \propto P(h|d) \left[\frac{P(h|d)}{P(h|-d)} \right]^\gamma. \quad (5)$$

which exactly captures Kahneman and Tversky’s (1983) definition of representativeness: “an attribute is representative of a class if it is very diagnostic; that is, the relative frequency of this attribute is much higher in that class than in a relevant reference class.” In this case, our model yields the BCGS (2016) definition of representativeness and diagnostic expectations (BGS 2018), which directly follow KT’s definition, as special cases of similarity based recall.

Our experiments test the predictions that follow from $\gamma > 0$ by varying the decoy data $-d$ across treatments. Our experiments then allow us to quantify the parameters of the model, and hence to assess Equation (4) quantitatively.

2.2 Experimental Design

We now describe our main experiments, and the model’s predictions for each of them.

Study 1. Participants are shown a sequence of 50 abstract images that vary along two dimensions, content and color. They are randomly assigned to one of two treatments: a condition in which they see 10 orange numbers, 15 blue numbers and 25 gray shapes (*gray* treatment); or a condition in which they see the same 10 orange numbers and 15 blue numbers, but also 25 blue words (the *blue* treatment). The memory databases are described below.

After observing the sequence of 50 images, participants are asked several questions. On the first screen, they are asked:

Table 1: Databases of Study 1

Image database in the <i>gray</i> treatment				
		Hypotheses:		
		orange	blue	gray
Data:	numbers	10	15	0
	decoy (shapes)	0	0	25

Image database in the <i>blue</i> treatment				
		Hypotheses:		
		orange	blue	gray
Data:	numbers	10	15	0
	decoy (words)	0	25	0

(Q1) “An image was randomly drawn from the images that were just shown to you. The chosen image showed a number. What is the likely color of the chosen image?”

On the following screen, participants are given the same scenario but then asked:

(Q2) “What is the probability the number is orange?”

Finally, on the next screen, participants are asked:

(Q3) “How many orange numbers were shown to you?”, and

(Q4) “How many blue numbers were shown to you?”

In both treatments, subjects assess the distribution of the colors of the numbers, but the decoy distribution varies across treatments. All our experiments build on this basic design, which extends the preliminary experiments in BCGS (2016).

In Q1 and Q2 subjects are cued with the data $d = \text{number}$ and are asked to assess the probability of different colors. In Q3 and Q4 they are asked to recall how many orange and blue numbers were shown to them. Our model implies that assessments of likelihood and recall go hand in hand, and are shaped by the critical contextual interference parameter γ in Equation (4). Formally, Appendix A shows that Equation (3) yields the following predictions.

Prediction 1. *If $\gamma > 0$, the blue treatment reduces the assessed likelihood that the randomly drawn number is blue in Q1 and Q2 relative to the gray treatment. If $\gamma = 0$ there is no treatment effect.*

and

Prediction 2. *If $\gamma > 0$, the blue treatment reduces the share of blue numbers recalled in Q3 and Q4 relative to the gray treatment. There is also a positive correlation at the individual level between the share of blue numbers recalled from Q3 and Q4 and the assessed likelihood of the number being blue in Q1 and Q2. If $\gamma = 0$ there is no treatment effect.*

If $\gamma > 0$, in the blue treatment the retrieval of blue words interferes with the retrieval of blue numbers from the data $d = \text{number}$. Thus, blue numbers are undersampled relative to the gray treatment, in which interference in the retrieval of numbers of any color is not at work. As a result, the probability of the less likely hypothesis – that a number is orange – is overestimated in the *blue* treatment. If instead $\gamma = 0$, the decoy distribution is irrelevant so there is no treatment effect.

Study 2. In this experiment, we manipulate the strength of contextual interference by creating five new variants of the decoy data (the distribution of numbers is always kept the same). We denote by *blue_k* a treatment in which we replace k blue words with orange words. We allow for six such treatments varying the frequency of orange words $k \in \{0, 1, 3, 6, 10, 22\}$ (where $k = 0$ coincides with the blue treatment in Study 1). The model generates the following prediction.

Prediction 3. *If $\gamma > 0$, the assessed probability that a random number is orange in Q1 and Q2 and the share of recalled orange numbers in Q3 and Q4 decrease with the number of orange words k in the decoy distribution. If $\gamma = 0$ there is no treatment effect.*

Prediction 3 is a more stringent test of the kernel of truth. According to Equation (4), the probability of retrieving orange given $d = \text{number}$ and thus the assessed likelihood of orange numbers both decrease in the share of orange words, because orange words interfere with the recall of orange numbers.

Study 3. In our final experiment, we test a key implication of the model: features of the task that shape perceptions of similarity affect memory retrieval and thus probability judgments, even if they are normatively irrelevant.

To explore this mechanism, we build on Study 1 where images are characterized by G (content) and T (color). We now let images vary also in a third attribute, specifically font size, $T' \in \{\textit{small}, \textit{large}\}$. The database for all participants is now:

All orange numbers are large and all blue numbers are small, while all words are blue and large.

We then run two treatments that vary only in the questions asked. In the Color treatment, we ask “What is the likely color of a randomly drawn number?” (Q1) and “What

Table 2: Database in Study 3

10 orange, large numbers
15 blue, small numbers
25 blue, large words

is the probability of a randomly drawn number being orange?" (Q2). In the Size treatment, we instead ask "What is the likely font size of a randomly drawn number?" and "What is the probability of a randomly drawn number being large?".

The point of this experiment is that the similarity of the data $d = \text{number}$ to a given color, say orange, is different from the similarity of the same data d to the corresponding size, namely large. This implies that asking subjects to assess color will produce different judgments than asking them to assess size, even though the two tasks are normatively identical.⁸ The model yields the following prediction.

Prediction 4. *If $\gamma > 0$, the share of recalled orange/large numbers and the assessment of the probability that a random number is orange/large is higher in the color treatment than in the size treatment. If $\gamma = 0$ there is no treatment effect.*

In the color treatment, contextual interference inhibits the recall of blue numbers due to the presence of blue words. In the size treatment, contextual interference inhibits the recall of large numbers due to the presence of large words. But because large numbers are orange, the size treatment inhibits the recall of orange numbers and their assessed likelihood relative to the color treatment. This effect of course only arises if $\gamma > 0$.

2.3 Discussion

Our model makes three distinctive predictions relative to the existing models of biased probabilistic assessments, and our experimental design allows us to test them.

First, our approach is based on selective recall. We measure recall directly in our experiment in Q3 and Q4. Some alternative approaches do not rely on basic cognitive functions, such as memory. Rather, they assume specific distortions without exploring their cognitive basis. For instance some papers assume that subjects compute $P(h|d)$ using $P(d|h)$ (Gigerenzer and Hoffrage 1995, Kahneman and Frederick 2002, Villejoubert

⁸Because in this experiment the distribution of images seen during the viewing stage is *identical* across treatments, there is no reason for a participant in one treatment to attend to or encode images differently from a participant in another treatment. This helps rule out the possibility that differential attention during encoding is driving our results. Although a precise distinction between selective encoding and selective recall is beyond the scope of this paper, it is useful to distinguish a process in which the instability of probabilistic assessments arises from the selectivity of retrieval from the memory database, rather than from a permanent distortion of the memory database itself, because of inattention at the encoding stage. In the current experiment, the first mechanism seems to be driving the results.

and Mandel 2002, Koehler et al. 2003, for a review see Nilsson et al. 2005), or that representative types provide strong confirmation of hypotheses (Tentori, Crupi, and Russo 2013). These distortions are tailored to specific inference problems, and lead to inaccurate predictions in many settings, including in our experiment.⁹

Second, relative to other models emphasizing memory limitations, our approach makes the distinctive "kernel of truth" prediction, whereby recall and assessments are context-dependent and exaggerate true correlations in the data (e.g, the association between the cue "number" and the type "orange"). This prediction is absent from other memory-based models. In sampling models (Sanborn and Chater 2016), subjects randomly sample the memory database, so the decoy distribution should not affect beliefs because it does not distort the sampling of numbers. Because in the Minerva DM model (Dougherty et al. 1999) retrieval is noisy, it generates imprecision but not systematic distortions depending on the decoy distribution.¹⁰ The same is true of exemplar models like PROBEX (Juslin and Persson 2002). In these models, subjects cued by "number" retrieve exemplars of different colors on the basis of a Euclidean similarity function that does not depend on the decoy distribution. As a consequence, these models cannot deliver Predictions 1, 3 and 4.

Evidence in line with Prediction 4 would offer particularly striking support to the role of contextual interference, which gives rise to the kernel of truth. Since in Study 3 size and color are perfectly correlated within the group of numbers, the cue should play no role for Bayesian subjects, nor should it play a role in sampling models (where numbers are correctly represented on average), in exemplar models (where large numbers and orange numbers are equally similar to the average number), or in prototype models (where the prototype is the most likely number, which is always blue).

⁹One intuition behind these models is that the inverse conditional $P(\text{number}|\text{orange})$ is more accessible to decision makers than the target one $P(\text{orange}|\text{number})$ (Kahneman and Frederick 2002). This might be the case when experimental subjects are given the former and are asked to estimate the latter, as in inference problems. This is less so in the general case where subjects retrieve a distribution from memory, where this approach leads to problems. For example, in the relative likelihood model of Villeboeuf and Mandel (2002) the assessed probability of orange numbers is over 50% as long as $P(\text{number}|\text{orange}) > P(\text{number}|\text{blue})$, even as the true share $P(\text{orange}|\text{number})$ goes to zero. A smoother formulation of the inverse conditional is a mechanical neglect of base rates, whereby the odds ratio $\frac{P(\text{orange}|\text{number})}{P(\text{blue}|\text{number})}$ is assessed as $\frac{P(\text{number}|\text{orange})}{P(n|b)} \left(\frac{P(\text{orange})}{P(\text{blue})} \right)^\phi$, with $\phi < 1$ (Bayes' rule corresponds to $\phi = 1$). This formulation predicts that if a prior is sufficiently strong, say $P(\text{orange}|\text{number})$ is high, then a signal that supports the prior (i.e. that the number is in fact orange) *reduces* the posterior probability assigned to that type. While our experiments do not cover this particular point, existing evidence generally indicates that individuals update their beliefs in the direction of their signals.

¹⁰Dougherty et al. (1999) argue that their model can also account for some manifestations of the representativeness heuristic. However, these effects are obtained by changing the model's formulation in ways that are unrelated with its core assumptions (namely, the authors assume that, when cued with data d and asked to assess hypothesis h , subjects spontaneously self-cue the hypothesis h and effectively assess $p(d|h)$, because the latter is assumed to be more accessible). Our model yields representativeness as a special case of a general process.

Finally, our approach to similarity is based on experienced statistical associations between items in memory. This occurs without the need of biases in social information transmission that emphasize stereotypical traits such as in Bhatia’s (2017) model of semantic associations. Our experiment rules out the role of such semantic associations by focusing on abstract objects.

3 Representativeness and Selective Recall

3.1 Study 1: Baseline Experiment

In Study 1 we compare probabilistic assessments in the *blue* treatment (10 orange numbers, 15 blue numbers, 25 blue words) and the *gray* treatment (10 orange numbers, 15 blue numbers, 25 gray shapes), as in Table 1.¹¹ According to Prediction 1, retrieval of blue numbers is interfered with in the *blue* treatment so that participants inflate the frequency of orange numbers relative to the *gray* treatment, $\tilde{P}(o|n)_{blue} > \tilde{P}(o|n)_{gray}$. Mapping Prediction 1 to our output measures, we test whether participants in the *blue* treatment: i) are more likely to choose orange as the likely color, ii) assign a higher probability to orange numbers, and iii) report a higher share of orange numbers compared to participants in the *gray* treatment, computed from answers to Q3 and Q4.

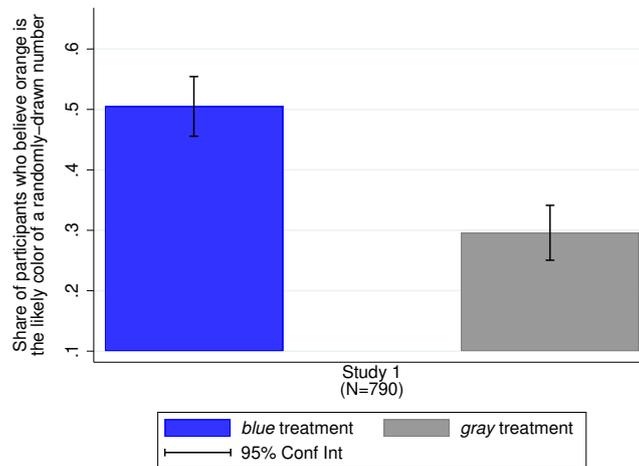


Figure 1: Share of participants who believe that the likely color of a randomly-drawn number is orange for the *blue* and *gray* treatments of Study 1.

The results are summarized in Figure 1 and Table 3 and show strong support for Prediction 1. Column (1) in Table 3 reports an OLS regression of a response dummy

¹¹This Study was conducted in the laboratory in Spring of 2018, in two waves with $N = 427$ (University of Cologne) and $N = 363$ (Bocconi University) respectively. Here we present the aggregated results. We report the procedural details and disaggregated results (Tables B.2 and B.3) in Appendix B.1.

(1 if “orange is likely”) on a treatment dummy (1 if *blue*) and a location dummy (1 if Milan), that amounts to comparing the average share of participants who said orange is likely in the *gray* treatment versus the *blue* treatment. As shown in Figure 1, that share increases 21.1pp from the *gray* treatment to the *blue* treatment (31.1% to 52.2%, significant at 1% level).¹² Column (2) shows the median probability of orange numbers is significantly higher in the *blue* treatment than in the *gray* treatment (significant at the 1% level).

Table 3: Regression estimates of treatment effects in Study 1 (lab with distraction)

	OLS: Y=1 if “orange is more likely”	0.5-Q-Reg: Y= Probability that a randomly-drawn number is orange	0.5-Q-Reg Y= Orange numbers recalled	0.5-Q-Reg Y= Blue numbers recalled	OLS: Y=1 if more orange numbers recalled	0.5-Q-Reg: Y= Share of orange to total numbers recalled
	(1)	(2)	(3)	(4)	(5)	(6)
1 if <i>blue</i>	.2111*** (.0342)	.07*** (.0137)	0 (.6078)	− 2*** (.5878)	.1667*** (.0329)	.0556*** (.0124)
Location dummy	yes	yes	yes	yes	yes	yes
Constant	.3113*** (.0284)	.43*** (.0144)	12*** (.5046)	15*** (.4880)	.2576*** (.0273)	.4444*** (.0114)
Observations	790	790	790	790	790	790
Adj./Ps. R ²	0.04	0.03	0.02	0.01	0.03	0.02

Notes: This table presents estimates of the treatment effect on several outcome measures relating to our model predictions in Study 1 with distraction. Columns (1) and (5) present OLS regression of response dummies (1 if “orange is likely” in Column (1) and 1 if “recalled more orange than blue numbers” in Column (5)) on a treatment dummy (1 if *blue*). Columns (2), (3), (4) and (6) present 0.5-quantile regressions of subjects’ stated probability that a random number is orange, number of recalled orange numbers, number of recalled blue numbers and share of recalled orange numbers, respectively, on the treatment dummy. In all regressions, a location dummy is included. Standard errors are in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Columns (3) and (4) show how the median quantity of recalled orange and blue numbers depends on the treatment.¹³ Responses in the *gray* treatment are quite accurate, as indicated by the constant term. In the *blue* treatment, participants retrieve fewer blue numbers, consistent with interference from blue words.

We next compare how data on recall (Q3 and Q4) matches assessments of likelihood (Q1 and Q2). First, we report in Column (5) the results of an OLS regression of a response dummy that takes value 1 if the participant reported more orange numbers in Q3 than blue numbers in Q4, which is an alternative measure of each participants’ belief about the likely color. Consistent with Column (1), there is a significant increase in the share of participants who recalled more orange than blue numbers in the *blue* treatment (25.8% vs 42.4%). The treatment dummy coefficients are close in Columns (1) and (5).

¹²Nonlinear Logit regressions yield similar results, Table B.1 in Appendix B.1.

¹³We report median, rather than mean, responses to the amount of numbers recalled because they are by construction less noisy. However, similar results hold for means, see Table B.1 in Appendix B.1.

In fact, answers to Q1, Q3 and Q4 are consistent (in the sense of answering “orange” in Q1 if and only if more orange numbers are recalled in Q3 than blue numbers in Q4) for 87.3% of participants.

Finally, we compute the ratio of orange numbers to total numbers recalled, which is an alternative measure of the (recalled) probability of orange numbers. Column (6) shows that, as predicted by the model, participants recalled on average a significantly higher share of orange numbers in the *blue* treatment (50% versus 44.4%). Again, the coefficients are close in Columns (2) and (6). In fact, the stated probability of orange numbers explains about 75% of the variability in the share of recalled orange numbers across individuals (see Column (1) of Table B.11 in Appendix B.3.1).

The results point to systematic distortions in the retrieval of information, leading to distorted beliefs in the direction predicted by the model of representativeness in Equation (3). The treatment effect represents a jump of over 50% in the frequency of assessing orange as the likely color (Columns 1 and 5), and a 13 to 16% increase in the estimate of the probability of a randomly drawn number being orange (Columns 2 and 6), so it is large both in absolute and in relative terms.¹⁴ This suggests that interference can account for why unlikely traits are accessible after specific cues, offering an explanation for some effects attributed to the representativeness heuristic (KT 1972). In particular, such distortions predictably arise in an environment of purely abstract objects with no pre-existing associations in memory.^{15,16}

3.2 Study 2: Varying Relative Likelihood

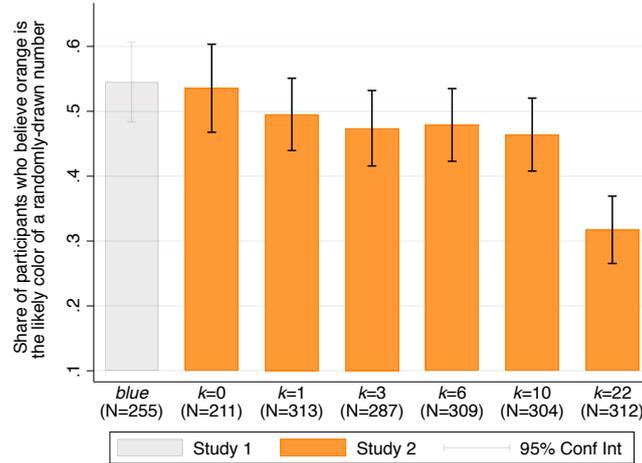
We next test whether, in line with Prediction 3, probabilistic assessments about the color distribution of numbers change as the frequency of orange words increases and interference for recall of orange numbers intensifies. As before, the target group is 25 Numbers

¹⁴The same experiment was also rerun in the lab and on MTurk without including question Q2 on the probability that a randomly drawn number is orange. The results are strikingly similar: the treatment effect represents a jump in the frequency of assessing orange as the likely color from 29% to 58% (lab) and 41% to 54% (MTurk), and an increase in the estimate of the share of orange numbers recalled from 45% to 48% (lab) and 47% to 50% (MTurk). See Tables B.4, B.5, B.6, and B.7 in Appendix B.1.1.

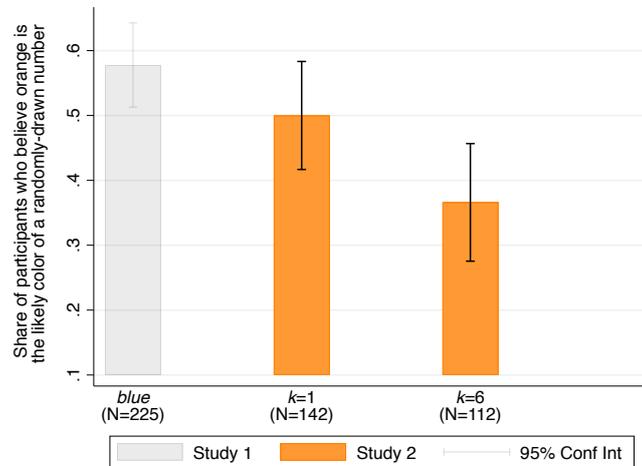
¹⁵One may ask whether the particular colors chosen impact our results. Research in psychology starting with Goldstein (1942) and Stone and English (1998) proposes that warmer colors, such as red and yellow, may induce outward focusing, while cooler colors may induce inward focusing (reservation), in part due to their different wavelengths. Elliot et al (2007) propose a more general framework in which the impact of specific colors varies by context and is a function of learned associations. It is unlikely that any pre-existing association between orange and numbers drives our results, particularly the across-treatment differences in the experiments where orange is not representative of numbers (see Study 2 and 3). We also ran a separate experiment, reported in Appendix C.3, that replaced color with font size as the types (i.e. the random variable T), and the results are unchanged (see Table C.5).

¹⁶A clear feature of the data is that the frequency of the “orange is likely” mistake is significant, around 35%, even in the gray treatments where our model predicts low distortions from representativeness. As we show when calibrating the model in Section 5, this distortions pins down the model parameter α , which creates probabilistic distortions even in the absence of contextual interference from a decoy group.

(10 orange and 15 blue). For the decoy group, we create six variants of the words distribution in the *blue* treatment, denoted $blue_k$, which are characterized by replacing k blue words with orange words. We conducted Study 2 on MTurk with $N = 1,738$ and $k = 0, 1, 3, 6, 10, 22$ and in the laboratory with $N = 254$ and $k = 1, 6$. Here we present the disaggregated results and discuss procedural details in Appendix B.2.



(a) Conducted with MTurk



(b) Conducted in Laboratory

Figure 2: Share of participants who believe that the color of a randomly-drawn number is most likely orange for the *blue* treatments with $k = 0, 1, 3, 6, 10, 22$.

As Figure 2 shows, the results are consistent with our predictions. For MTurk experiments (top panel), the share of participants who answered “orange is likely” decreases from 54.1% from the baseline treatment with $k = 0$ to 31% in the variant with most orange words, $k = 22$. The figure suggests a decline between $k = 0$ and $k > 0$ that becomes particularly strong from $k = 10$ to $k = 22$. The decline is statistically significant at the 5% level between $k = 0$ and $k \geq 6$ and at the 1% level between $k \leq 10$ and

$k = 22$.¹⁷

Similar results hold in laboratory experiments, shown on the lower panel. Increasing the number of orange words from 1 to 6 reduces the share of participants stating that “orange is likely” by 21.2pp (significant at the 5% level). This again points to a general trend, as can be seen by comparing these findings with the results from Study 1.¹⁸

Table 4 summarizes the pattern described above for the experiments run in the lab. Column (1) reports an OLS regression of a response dummy (1 if “orange is likely”) on the actual amount of orange words participants are exposed to. The significant negative coefficient implies that the share of participants who believed that “orange is likely” decreases by 2.7pp per orange word added. Similarly, the median assessed likelihood of orange numbers drop by 2pp per orange word added (Column 2). Turning to data on recall, the share of participants who recalled seeing more orange numbers as well as the share of orange to total numbers recalled declines in the amount of orange words, Column (5) and Column(6), respectively.¹⁹

Again, results obtained on MTurk are quantitatively very similar. The share of participants who say “orange is likely” drops by 1pp per orange word added. Turning to data on recall, as orange words are added, participants recall more blue numbers and are less likely to recall seeing more orange numbers (see Table B.10 in Appendix B.2).

Study 2 provides evidence that the relative frequency of types shapes the magnitude of belief distortions. In line with Prediction 3, as the number of orange words increases, participants are less likely to recall a randomly-drawn number as orange, despite the fact that the number of orange numbers is held constant across these variants.

Studies 1 and 2 include waves of our experiment in the Lab with 1044 participants (further robustness runs described in Appendix B.1.1 raise these figures to 8 waves of our experiment and a total of 4706 participants). The evidence supports our predictions consistently across all waves of our experiment.

¹⁷Increasing the number of orange words from 0 to 1, 3, 6, 10, and 22 reduces the share of MTurk participants stating that “orange is likely” by 5.3pp, 7.4pp, 7.0pp, 8.4pp, and 23.1pp, respectively. While the difference between 0 orange words and 1 orange words is not statistically significant, the remaining ones are at least marginally significant when compared to the $blue_{k=0}$ treatment, with respective OLS p -values of 0.070, 0.085, 0.037, and < 0.001 . Pair-wise differences in the assessed probability that a randomly selected number is more likely to be orange comparing across only those treatments with 1, 3, 6, and 10 orange words are not large and not statistically significant from zero. However, the assessed probability that a random number is more likely to be orange is greater in the treatments with 1, 3, 6, or 10 orange words than in the treatment with 22 orange words. Pair-wise tests yield OLS p -values below 0.001 in each of these cases.

¹⁸Study 1’s $blue$ treatment is equivalent to $k = 0$. The $gray$ treatment is not directly comparable, but to the extent that no number color is particularly representative in that treatment it is similar to $k = 10$.

¹⁹Table 4 shows no effect on the median number of recalled blue or orange numbers from the addition of five orange words. Some evidence suggests that this is due to power limitation of this exercise: i) there is a significant negative effect on the mean (see Appendix B.2), and ii) there is also a significant effect in the median in the MTurk experiment, where a larger range of orange words was introduced.

Table 4: Regression estimates of treatment effects in Study 2 (lab only)

	OLS: Y=1 if “orange is likely”	0.5-Q-Reg: Y= Probability that a ran- dom number is orange	0.5-Q-Reg: Y= Orange numbers recalled	0.5-Q-Reg: Y= Blue numbers recalled	OLS: Y= 1 if more orange numbers recalled	0.5-Q-Reg: Y= Share of orange to total num- bers recalled
	(1)	(2)	(3)	(4)	(5)	(6)
k (number of orange words)	-.0268** (.0125)	-.02*** (.0071)	0 (.1466)	0 (.2108)	-.0164 (.0121)	-.0123** (.0062)
Constant	.5268*** (.0506)	.5200*** (.0288)	10*** (.4867)	15*** (.5943)	.4107*** (.0491)	.4738*** (.0251)
Observations	254	254	254	254	254	254
Adj./Ps. R^2	0.02	0.02	0.00	0.00	0.01	0.01

Notes: This table presents estimates of the treatment effect on several outcome measures relating to our model predictions in Study 2. We report on our results from the lab here, because we elicited participants’ expected probability that a random number is orange not via MTurk. Columns (1) and (5) present OLS regression of response dummies (1 if “orange is likely” in Column (1) and 1 if “recalled more orange than blue numbers” in Column (3)) on k , the number of orange words. Columns (2), (3), (4) and (6) present 0.5-quantile regressions of subjects’ stated probability that a random number is orange, number of recalled orange numbers, number of recalled blue numbers and share of recalled oranges numbers, respectively, on $k = 6$. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

4 Study 3: Modulating Recall through Cues

In the final experiment, we attempt to better isolate the role of similarity-based recall in driving our results by testing Prediction 4: cueing assessments of a target group along different dimensions can trigger different patterns of interference across groups and generate different probabilistic assessments.

We implement the distribution of images described in Table 2, where all blue numbers are small while all other items (orange numbers and blue words) are large. According to Prediction 4, assessment of a number’s size entails interference with recall of large (and orange) numbers, while assessment of a number’s color entails as before interference with recall of blue (and small) numbers. As a result, the estimated likelihood of large, orange numbers is larger in the second case.

The results support this prediction. Participants assess “orange is likely” in the Color Cue treatment significantly more often than “large is likely” in the Size Cue treatment (40% versus 17%, significant at the 1% level). Column (1) of Table 5 shows the result of regressing a question-response dummy (equal 1 if “orange is likely” in Color Cue or “large is likely” in Size Cue) on a treatment dummy (equal 1 if Color Cue, equal 0 if Size Cue), while controlling for when the treatments were conducted.

Table 5: Regression estimates of treatment effects in Study 3

	OLS: Y=1 if “orange OR large is likely”	OLS: Y= 1 if more orange OR large numbers recalled	0.5-Q-Reg: Y= Share of orange OR large to total numbers recalled	0.5-Q-Reg: Y= Probability that a randomly-drawn number is orange OR large
	(1)	(2)	(3)	(4)
1 if color cue	.2296*** (.0343)	.1168*** (.0341)	.04** (.0178)	.05** (.0218)
Wave dummy	yes	yes	yes	yes
Constant	.1808*** (.0293)	.1953*** (.0291)	.41*** (.0152)	.35*** (.0186)
Observations	647	647	647	647
Adj./Ps. R^2	0.06	0.02	0.01	0.01

Notes: This table presents estimates of the treatment effect on several outcome measures relating to our model predictions. Columns (1) and (2) present an OLS regression of response dummies (1 if “orange is likely” in treatment color cue and 1 if “large is more likely” in treatment font size cue for Column (1) and 1 if “more orange numbers are recalled than blue numbers” in treatment color cue and 1 if “more large numbers are recalled than small numbers” in treatment font size cue for Column (2)) on a treatment dummy. Columns (3) and (4) present 0.5-quantile regressions of subjects’ share of recalled orange (treatment *color cue*) OR large (treatment *font size cue*) numbers over total recalled numbers and stated probability that a random number is orange (treatment *color cue*) OR large (treatment *font size cue*), respectively, on the treatment dummy. In all regressions, a wave dummy is included. Standard errors are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

As in Tables 3 and 4, we use participants’ estimates of how many images of each type they saw to address more directly their retrieval of numbers. First, we compare the share of participants who recalled seeing more orange than blue numbers in Color Cue with the share of participants who recalled seeing more large than small numbers in Size Cue. We find, analogously to our main result, that a significantly greater share of participants recalled more orange than blue numbers than more large than small numbers; see Column (2) of Table 5. Second, we check the average share of recalled large orange numbers and find, consistent with the account of interference-based distorted recall, that on average participants recalled a significantly higher share of large orange numbers under the color cue, see Column (3) of Table 5.

Finally, we examine the results from direct probability estimates in Q4. When asked to predict the probability that a randomly-drawn number is orange/large, participants on average state a significantly higher probability that a random number is orange in the Color Cue treatment than that a random number is large in the Size Cue treatment; see Column (2) of Table 5. These findings suggest that participants’ retrieval of blue and orange numbers from their image database was differentially distorted depending on

whether color or font size was cued in the question stage.

5 Model Calibration and Assessment

The previous sections document that elicited beliefs about a given group, numbers, are shaped by its comparison to a reference group, in line with the predictions of our model. In fact, there is significant and systematic variation across treatments in subjects' average estimation of the likelihood $\tilde{P}(o|n)$, ranging from 0.36 to 0.499. To assess how well the model can explain the observed variation in beliefs, and to quantitatively estimate the role of interference in the model, we conduct a calibration exercise.

Beliefs about numbers can be summarized in the elicited odds $\frac{\tilde{P}(o|n)}{\tilde{P}(b|n)}$ of orange to blue numbers. Our calibration is targeted to match this moment as it varies across treatments. Starting from Equation (3), we obtain:

$$\frac{\tilde{P}(o|n)}{\tilde{P}(b|n)} = e^{S(n,o)-S(n,b)}.$$

Assuming, as in Section 2, that $\beta = 0$ and using the same logarithmic functional form $f(x) = \ln(c + x)$, the predicted assessed odds of orange numbers are given by:

$$\frac{\tilde{P}(o|n)}{\tilde{P}(b|n)} = \frac{(c + P(o, n))^{\alpha-\gamma} \cdot \left(\frac{c+P(o,n)}{c+P(o,-d)}\right)^\gamma}{(c + P(b, n))^{\alpha-\gamma} \cdot \left(\frac{c+P(b,n)}{c+P(b,-d)}\right)^\gamma} \quad (6)$$

where $P(o, n)$ and $P(b, n)$ are the frequency of orange and blue numbers among all images shown to participants, while $P(o, -d)$, $P(b, -d)$ are the frequency of orange and blue decoy images among all images shown to participants.

Table 6

Study	Treatment $-d$	Experimental Odds
1 (lab)	gray shapes	0.78
1 (lab)	words $k = 0$	0.99
2 (lab)	words $k = 1$	0.93

Notes: Target moments.

We next pin down the three parameters α , γ and c by matching the odds (6) with the odds computed from the average assessment across subjects of the likelihood of orange numbers. We match three predicted odds to their empirical counterparts in three

different conditions. Table 6 shows the treatments used for calibration and the corresponding empirical target odds. The first two moments come from the *blue* and *gray* conditions of our baseline Study 1 in the Lab. The third moment comes from Study 2, and in particular from the treatment with one orange word ($k = 1$). Across these three conditions, in Equation (6) the target distribution of numbers is unchanged, while the decoy distribution changes.²⁰

We invert the resulting system of three Equations (6) using standard numerical methods, and obtain:

$$\begin{array}{ccc} \alpha^* & \gamma^* & c^* \\ \hline 0.66 & 0.07 & 0.01 \\ \hline \end{array}$$

Our numerical solution achieves a good fit of the empirical odds (least squared error of 9×10^{-4} , depicted in Figure 4) which suggests that the estimates are very close to the exact solution of the system. Parameter γ is tightly estimated, as shown in Figure 3, indicating an important role for interference across groups in shaping beliefs. According to the model, replacing gray shapes with blue words generates a 7pp increase in the perceived probability of orange numbers, from 0.43 to 0.50.

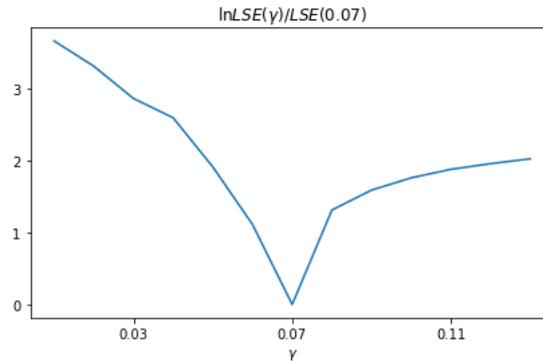


Figure 3: Model fit. For each value of γ , we compute the least square error over parameters α, c , namely $LSE(\gamma) = \min_{\alpha, c} \sum_i \left(\frac{\tilde{P}(on)_i}{\tilde{P}(b|n)_i} - \frac{P(on)_i}{P(b|n)_i} \right)^2$ where i indexes the three target moments in Table 6. $LSE(\gamma)$ attains a minimum at $\gamma^* = 0.07$. The Figure plots the ratio $\frac{\ln LSE(\gamma)}{\ln LSE(\gamma^*)}$ as a function of γ .

Using these estimates, we next examine the model's performance out of sample in: i) assessments of likelihood in treatments not targeted by the calibration, ii) assessments of modal color across Studies 1, 2 and 3, iii) recall of orange and blue numbers across Studies 1, 2 and 3, and iv) recall of orange and blue numbers in supplementary experiments where the number of decoys is varied.

²⁰The experimental data includes six moments on probabilistic assessments of the probability of orange numbers, which can be mapped directly to the model. The three chosen target moments come from the three most basic experiments. We then use the remaining moments for out-of-sample tests of model performance.

We start by computing the model’s predictions for assessments $\tilde{P}(o|n)$ in the remaining experiments where we elicited likelihood assessments, namely in Lab Study 2, $g = \text{words}(k = 6)$ and in Lab Study 3, both the size and color treatments. Figure 4 shows the results (blue dots), as well as the fit with the target moments (blue circles). To facilitate the assessment of model performance, Figure 4 also shows the predictions of the calibrated model where interference is shut down ($\gamma = 0$, dotted line) as well as the 45° line that would obtain if the model would perfectly match the evidence. The calibrated model reproduces the large observed variation in the assessed probability of orange numbers, and matches fairly well the out of sample moments, in particular the impact of the size cue which leads to a large drop in $\tilde{P}(o|n)$ to a value much lower than in the target treatments. The calibrated model thus significantly outperforms models that lack contextual interference, $\gamma = 0$, which would predict no variation in $\tilde{P}(o|n)$ (the horizontal dotted line).

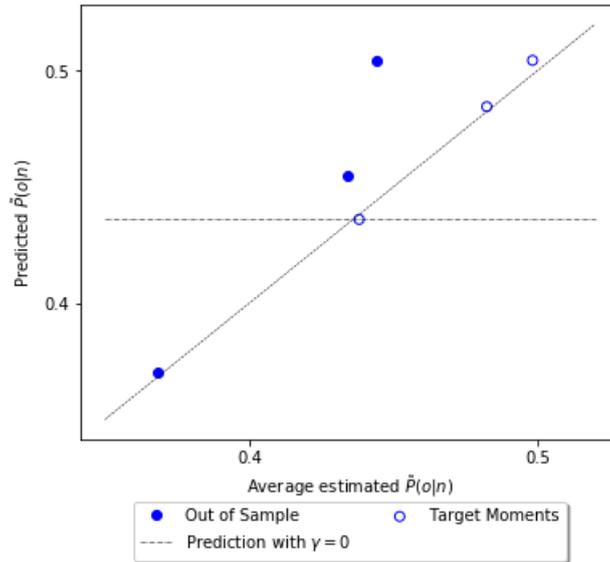


Figure 4: Model performance I: assessments of likelihood, data $\hat{P}(o|n)$ vs model predictions $\tilde{P}(o|n)$. The graph also plots the 45° line, as well as the prediction of the calibrated model where contextual interference is shut down, $\gamma = 0$.

We next consider subjects’ assessments of which color is modal, which are closely related to likelihood assessments. In every experiment, some subjects say orange and some say blue, reflecting disagreement that is also observed in their assessments of likelihood. To assess model performance in this context, it is not sufficient to use the average assessed probability of orange numbers (which is weakly below 0.5 in every experiment). For the purpose of obtaining a distribution of assessments across subjects, we assume that they sample numbers from memory from a binomial distribution of colors with a probability of orange given by the model predicted likelihood $\tilde{P}(o|n)$. Specifi-

cally, we assume they sample $N = 25$ numbers, which is the true number of numbers and also coincides with the modal number of numbers recalled. Naturally, averaging the likelihood of orange numbers across subjects again yields $\tilde{P}(o|n)$. But this process also yields a share of subjects that state that orange is the modal color, which is given by $1 - cdf(N/2, N, \tilde{P}(o|n))$, where cdf stands for the cumulative density function of the binomial distribution.²¹ Figure 5 plots the resulting predicted shares against the empirical shares for our studies, which were obtained both in the Lab and on Mturk.²²

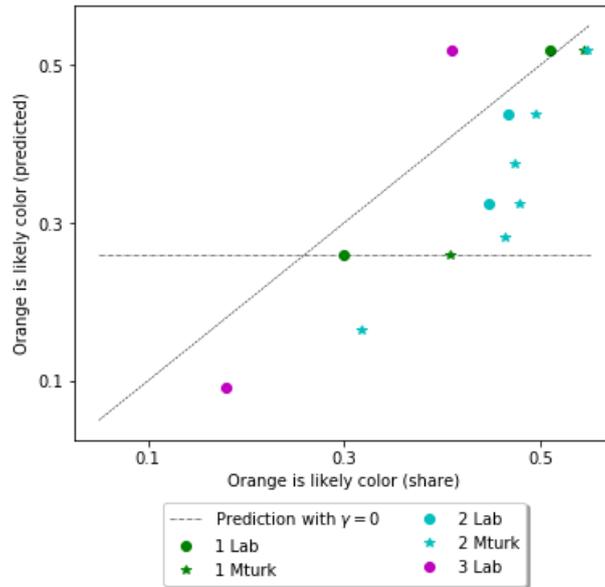


Figure 5: Model performance II: assessments of modal color; data vs model predictions. The graph also plots the 45° line, as well as the prediction of the calibrated model where contextual interference is shut down, $\gamma = 0$.

Again, the model captures quantitatively the large variation in the share of subjects reporting that orange is the modal color, which varies by a factor of nearly three (from 0.18 to 0.52) – a pattern that is absent in models that lack contextual interference (horizontal dotted line). The model tends to underpredict this share, particularly relative to the MTurk data, which may be due to the fact that with binomial sampling mistakes about the mode become very unlikely as $\tilde{P}(o|n)$ drops below 0.5, which may be too strict (the assumption that this sampling process is unbiased may not hold in reality, since thinking about an orange number may cue further retrieval of orange numbers).

We next go back to the data on recall. From each subject’s recalled number of orange and blue numbers we construct a recalled share of orange numbers. We then compare the

²¹We do not pursue further the analysis of the resulting distribution of likelihood assessments, because the results rely heavily on the assumed sampling process which is not an integral part of our model.

²²Recall that Lab runs of Study 2 included treatments $k = 1, 6$ while MTurk runs included treatments $k = 1, 3, 6, 10, 22$. Moreover, Study 3 was only run in the Lab.

recalled share in each treatment $-d$, averaged across subjects, to the model’s predicted probability $\tilde{P}(o|n)$ in that treatment. Figure 6 shows the results.

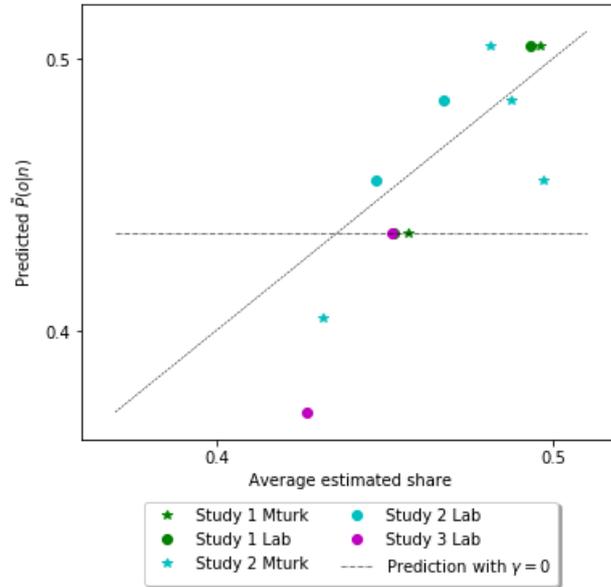


Figure 6: Model performance III: recall of orange numbers vs model predicted $\tilde{P}(o|n)_{g'}$. The graph also plots the 45° line, as well as the prediction of the calibrated model where contextual interference is shut down, $\gamma = 0$.

The model again matches well the recalled shares, particularly in Studies 1 and 2. This is perhaps not surprising given that the correlation between recall and assessed probabilities is very high in these studies (see Table B.11 in Appendix B.3.1). This high correlation is a non-trivial result consistent with the model’s predictions that assessments of probability distributions are shaped by recall of group elements.²³

Finally, we also assess the performance of the model relative to Study 1.A, described in Appendix C.1, which explores a different source of variation in the similarity between colors and numbers: group size. In this experiment, as before, we keep constant the set of numbers (15 blue, 10 orange) as well as the color distribution among words (all blue), but vary the number of words from 5 to 50, 75 and 125. Consider the effect of adding more blue words. According to the similarity function (2), this: i) reduces the similarity of any color to numbers, because there numbers are a smaller share of all objects, and ii) increases interference with the recall of blue numbers. Thus, adding blue words increases the (relative) similarity of orange to numbers and the assessed probability $\tilde{P}(o|n)$. Figure 7 (green dots) shows that the model’s prediction holds in the data. Moreover, the

²³One outlier is Study 3’s size treatment, where recalled share is much higher than assessed likelihood. This may be due to the fact that the differential cueing (size vs color) in that experiment was stronger for the elicitation of likelihood (which occurred first) than for the recall task, in which subjects were asked to recall numbers both in terms of size and of color.

calibrated model matches well the quantitative variation in assessed probability of orange numbers as the number of words changes, even though the experiment highlights a source of variation not used in the calibration.

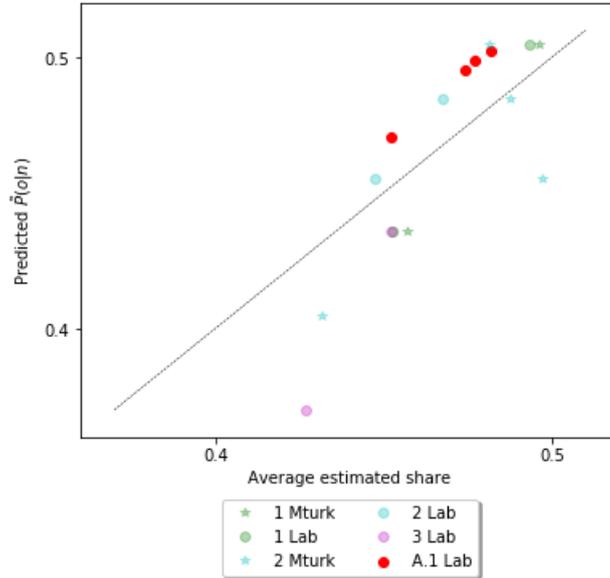


Figure 7: Model performance IV: recall of orange numbers vs model predicted $\tilde{P}(o|n)_{g'}$, Study A.1. The graph also plots the 45° line.

We conclude this Section with a discussion of the other model parameters. As shown in Equation (2), a value of α below 1 means that the similarity between a hypothesis and data does not respond one for one to their overlap: adding more blue numbers does not make the set of blue objects substantially more similar to the set of numbers. As a consequence, memory-based beliefs underreact to the variation in likelihood of hypotheses given data, and tilt their odds towards 1. This underreaction leads to an overestimation of rare types, at the cost of underestimating frequent types. In our experiments, this contributes to an overestimation of orange numbers, which however is small in magnitude and constant across treatments.

Finally, parameter c captures a baseline probability of any hypothesis coming to mind. This parameter plays an important role in smoothing the effects of extreme representativeness, which arise in the similarity expression (2) when we use the logarithmic functional form for $f(\cdot)$.

In sum, our model of selective memory offers a unified quantitative account of experimental measurements on recall and probability assessments using the new mechanism of contextual interference.

6 Discussion

This paper explores the link between intuitive thinking, as exemplified by the representativeness heuristic, and memory. The evidence takes the form of a systematic instability of probabilistic assessments: numbers are recalled as being more likely to be orange than blue when presented together with blue words than with orange words. This evidence shows that the probabilistic assessment of a hypothesis (object is orange) given data (object is a number) involves selective retrieval of hypotheses from memory, and that this retrieval is subject to contextual interference from other data. The role of cue-based recall and of contextual interference is most clearly illustrated by the fact that, keeping experience constant, probabilistic judgments about given data are dramatically altered depending on what comparisons are triggered by the cue (Study 3).

The crucial theoretical ingredient is that recall is driven by a measure of similarity that captures the statistical correlation between hypotheses and data, including contextual interference. This measure, which is related to Tversky's (1977) contrast model, nests representativeness (GS 2010, BCGS 2016), and captures the observed alignment between similarity judgments and probability assessments (Kahneman and Tversky 1983). The notion that similarity dependent on context, originally motivated by intuitive judgments of similarity (Tversky 1977), seems central also to account for similarity-based recall tasks including probabilistic assessments. The model highlights how, through contextual interference, such assessments are determined from objectively experienced statistical associations.

As described in the Introduction, the patterns of interference uncovered by our experiment seem to be at the heart of many examples of base-rate neglect, as well as the conjunction fallacy (Gennaioli Shleifer 2010).²⁴ Interference is also consistent with the Cognitive Psychology approach to stereotypes, as described by Hilton and Hippel (1996): "stereotypes are selective [...] in that they are localized around group features that are the most distinctive, that provide the greatest differentiation between groups." Here again, contextual interference seems essential.

Some stereotypical beliefs can be amplified by exposure to biased sources of information that confirm intuitive beliefs, such as natural language (Bhatia 2017). Such semantic associations may be unavoidable in driving beliefs about groups that are particularly salient in real world situations. In these cases, we view the two approaches as complementary. Interference in retrieval may help predict which specific features of groups

²⁴Consider the Linda problem: given Linda's background, a variety of professional outcomes are possible. Retrieval of the bank teller outcome is dampened because the average bank teller is more strongly associated with people of different, perhaps less rebellious, backgrounds. Subjects thus underestimate the probability that Linda is a bank teller for the same reason they may overestimate the probability that she is a social worker. In comparison to the generic bank teller, the more specific feminist bank teller outcome is much more representative of Linda, so (between subjects) a higher likelihood is assigned to it.

will form a stereotype. Natural language and rehearsal may disproportionately sample a group's representative types, reinforcing the belief originating in selective recall. In this way, selective recall and semantic similarity are likely to be complementary forces in creating real world judgments.

Overall, our results suggest that memory shapes the accessibility of thoughts that drive intuitive thinking about past experiences and can be an important determinant of probabilistic beliefs (Kahneman 2003). This suggests that other distortions in probabilistic judgment might be understood by integrating other memory mechanisms – not only contextual interference – into settings where people form probabilistic assessments by retrieving information from memory on the basis of cues. For instance, it has been suggested that anchoring biases might be due to selective retrieval based on similarity with the anchor (Strack and Mussweiler 1997). Related, it is possible that providing the inverse conditional probability of data given hypothesis anchors judgment of the probability of hypothesis given data (as in the taxicab problem) precisely because the former is very accessible to memory. The availability heuristic, whereby probabilistic judgments reflect “the ease with instances or occurrences can be brought to mind” (KT 1974, page 1127), explicitly relies on recall. Finally, semantic memory might be at play when making probabilistic judgments in non-experiential settings, where the likelihood of an event may be “judged by the degree to which it reflects the salient features of the process by which it is generated” (KT 1972, page 431). The notion of memory here may be to retrieve specific processes in light of available cues, perhaps along the lines advocated by Schacter (2007): “a crucial function of memory is to make information available for the simulation of future events” (p. 659). Extending our framework to incorporate these issues may provide a unified explanation of judgment biases and intuitive thinking.

References

- [1] John Anderson. Retrieval of propositional information from long-term memory. *Cognitive Psychology*, 6(4):451–474, 1974.
- [2] John Anderson and Lynne Reder. The fan effect: New results and new theories. *Journal of Experimental Psychology: General*, 128(2):186–197, 1999.
- [3] Michael Anderson and Barbara Spellman. On the status of inhibitory mechanisms in cognition: Memory retrieval as a model case. *Psychological Review*, 102(1):68–100, 1995.
- [4] Daniel Benjamin. Errors in probabilistic reasoning and judgment biases. In Douglas Bernheim, Stefano DellaVigna, and David Laibson, editors, *Handbook of Behavioral Economics: Foundations and Applications 2*, pages 69–186. Elsevier, 2019.
- [5] Sudeep Bhatia. Associative judgment and vector space semantics. *Psychological Review*, 124(1):1–20, 2017.
- [6] Pedro Bordalo, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer. Stereotypes. *Quarterly Journal of Economics*, 131(4):1753–1794, 2016.
- [7] Pedro Bordalo, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer. Beliefs about gender. *American Economic Review*, 109(3):739–73, 2019.
- [8] Pedro Bordalo, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer. Diagnostic expectations and stock returns. *Journal of Finance*, 74(6):2839–2874, 2019.
- [9] Pedro Bordalo, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. Overreaction in macroeconomic expectations. Technical report, Harvard University, 2018.
- [10] Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Diagnostic expectations and credit cycles. *Journal of Finance*, 73(1):199–227, 2018.
- [11] Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Memory, attention, and choice. Technical report, Harvard University, 2018.
- [12] Aaron Bornstein and Kenneth Norman. Reinstated episodic context guides sampling-based decisions for reward. *Nature neuroscience*, 20(7):997, 2017.
- [13] Alberto Cavallo, Guillermo Cruces, and Ricardo Perez-Truglia. Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. *American Economic Journal: Macroeconomics*, 9(3):1–35, 2017.

- [14] Michael Dougherty, Charles Gettys, and Eve Ogden. Minerva-dm: A memory processes model for judgments of likelihood. *Psychological Review*, 106(1):180–209, 1999.
- [15] Andrew Elliot, Markus Maier, Arlen Moller, Ron Friedman, and Jörg Meinhardt. Color and psychological functioning: The effect of red on performance attainment. *Journal of Experimental Psychology: General*, 136(1):154–168, 2007.
- [16] Benjamin Enke, Frederik Schwerter, and Florian Zimmermann. Associative memory and belief formation. Technical report, Harvard University, 2019.
- [17] Nicola Gennaioli and Andrei Shleifer. What comes to mind. *Quarterly Journal of Economics*, 125(4):1399–1433, 2010.
- [18] Gerd Gigerenzer and Ulrich Hoffrage. How to improve bayesian reasoning without instruction: frequency formats. *Psychological Review*, 102(4):684–704, 1995.
- [19] Kurt Goldstein. Some experimental observations concerning the influence of colors on the function of the organism. *Occupational Therapy*, 21:147–151, 1942.
- [20] James Hilton and William Von Hippel. Stereotypes. *Annual Review of Psychology*, 47(1):237–271, 1996.
- [21] John Jenkins and Karl Dallenbach. Obliviscence during sleep and waking. *The American Journal of Psychology*, 35(4):605–612, 1924.
- [22] Peter Juslin and Magnus Persson. Probabilities from exemplars (probex): A “lazy” algorithm for probabilistic inference from generic knowledge. *Cognitive Science*, 26(5):563–607, 2002.
- [23] Michael Kahana. *Foundations of Human Memory*. OUP USA, 2012.
- [24] Daniel Kahneman. Maps of bounded rationality: psychology for behavioral economics. *American Economic Review*, 93(5):1449–1475, 2003.
- [25] Daniel Kahneman and Shane Frederick. Representativeness revisited: Attribute substitution in intuitive judgment. In Thomas Gilovich, Dale Griffin, and Daniel Kahneman, editors, *Heuristics and biases: The psychology of intuitive judgment*, pages 49–81. Cambridge University Press, 2002.
- [26] Daniel Kahneman and Amos Tversky. Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3(3):430–454, 1972.

- [27] Daniel Kahneman and Amos Tversky. Choices, values, and frames. *American Psychologist*, 39(4):341–350, 1984.
- [28] Charles Kemp, Aaron Bernstein, and Joshua Tenenbaum. A generative theory of similarity. In *Proceedings of the 27th annual conference of the cognitive science society*, pages 1132–1137. Citeseer, 2005.
- [29] Geoffrey Keppel. Retroactive and proactive inhibition. In Theodore R Dixon and David L Horton, editors, *Verbal Behavior and General Behavior Theory*, pages 172–213. Prentice Hall, 1968.
- [30] Derek Koehler, Chris White, and Ray Grondin. An evidential support accumulation model of subjective probability. *Cognitive Psychology*, 46(2):152–197, 2003.
- [31] Ulrike Malmendier and Stefan Nagel. Depression babies: Do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics*, 126(1):373–416, 2011.
- [32] John McGeoch. Forgetting and the law of disuse. *Psychological Review*, 39(4):352–370, 1932.
- [33] Don Moore and Paul Healy. The trouble with overconfidence. *Psychological review*, 115(2):502, 2008.
- [34] Håkan Nilsson, Henrik Olsson, and Peter Juslin. The cognitive substrate of subjective probability. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(4):600, 2005.
- [35] Fred Paas, Alexander Renkl, and John Sweller. Cognitive load theory and instructional design: Recent developments. *Educational Psychologist*, 38(1):1–4, 2003.
- [36] Adam Sanborn and Nick Chater. Bayesian brains without probabilities. *Trends in Cognitive Sciences*, 20(12):883–893, 2016.
- [37] Nancy Stone and Anthony English. Task type, posters, and workspace color on mood, satisfaction, and performance. *Journal of Environmental Psychology*, 18(2):175–185, 1998.
- [38] Joshua B Tenenbaum and Thomas Griffiths. The rational basis of representativeness. In *Proceedings of the 23rd Annual Meeting of the Cognitive Science Society*, pages 1036–1041, 2001.
- [39] Katya Tentori, Vincenzo Crupi, and Selena Russo. On the determinants of the conjunction fallacy: Probability versus inductive confirmation. *Journal of Experimental Psychology: General*, 142(1):235–255, 2013.

- [40] Amos Tversky. Features of similarity. *Psychological Review*, 84(4):327–352, 1977.
- [41] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.
- [42] Amos Tversky and Daniel Kahneman. Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90(4):293–315, 1983.
- [43] Benton Underwood. Interference and forgetting. *Psychological Review*, 64(1):49–60, 1957.
- [44] Gaëlle Villejoubert and David Mandel. The inverse fallacy: An account of deviations from bayes’s theorem and the additivity principle. *Memory & Cognition*, 30(2):171–178, 2002.
- [45] Jessica Wachter and Michael Kahana. A retrieved-context theory of financial decisions. Technical report, University of Pennsylvania, 2019.
- [46] Paul Whitely. The dependence of learning and recall upon prior intellectual activities. *Journal of Experimental Psychology*, 10(6):489–508, 1927.

Memory and Representativeness

Bordalo, Coffman, Gennaioli, Schwerter, Shleifer

Supplementary Materials

A Derivation of the Predictions

To derive the predictions, we start from Expression (3) for the average assessed probability of hypothesis $h = \text{orange}$ given data $d = \text{numbers}$, and obtain:

$$\tilde{P}(h|d) = \frac{e^{S(d,h)}}{e^{S(d,h)} + e^{S(d,-h)}} = \frac{1}{1 + e^{S(d,-h)-S(d,h)}}$$

Thus, $\tilde{P}(h|d)$ is a monotonically increasing function of

$$S(d,h) - S(d,-h) = \alpha [f(P(d,h)) - f(P(d,-h))] - \gamma [f(P(-d,h)) - f(P(-d,-h))]$$

It follows that:

$$\frac{d\tilde{P}(h|d)}{dP(h|d)} \propto \alpha [f'(P(h|d)P(d)) + f'((1 - P(h|d))P(d))]P(d)$$

which is positive if and only if $\alpha > 0$, because $f()$ is an increasing function. Similarly, we find:

$$\frac{d\tilde{P}(h|d)}{dP(h|-d)} \propto -\gamma [f'(P(h|-d)P(-d)) + f'((1 - P(h|-d))P(-d))]P(-d)$$

which is negative if and only if $\gamma > 0$.

Predictions 1 and 3 then follow from the latter comparative static. For Prediction 1, $P(\text{blue}|\text{words}) > P(\text{blue}|\text{shapes}) = 0$ so that

$$\tilde{P}(\text{blue}|\text{numbers})_{d=\text{words}} < \tilde{P}(\text{blue}|\text{numbers})_{d=\text{shapes}}$$

For Prediction 3, $P(\text{blue}|\text{words})_k$ decreases in the number k of orange words, so that $\tilde{P}(\text{blue}|\text{numbers})_k$ increases in that number.

Prediction 2 reflects the interpretation of Equation (3) as the probability with which hypothesis h comes to mind when cued with data d .

Finally, Prediction 4 reflects the role of the cue. Formally, denote the hypothesis space by $H \times G$, so that a complete hypothesis is a pair (h, g) . In our setting, given data d or $-d$, each $h \in H$ only occurs with a single $g \in G$ and vice versa. When the cue is H (e.g. color), denote the probability of hypothesis h with data $-d$ as $P(h(g)|-d)$.

When the cue is G (e.g. size), denote the probability of hypothesis g with data $-d$ as $P(g(h)|-d)$. Prediction 4 then follows from the fact that $P(\text{blue}(\text{large})|\text{words}) > P(\text{small}(\text{orange})|\text{words}) = 0$.

B Methods, Procedures, and Further Results

B.1 Study 1

Procedural details and data collection

In this section, we provide further details about Study 1. We conducted two waves of the *blue* and *gray* treatments. These experiments were conducted in May of 2018 in the laboratories of the University of Cologne (N=427) and at Bocconi University (N=363). The entire lab experiment took 10 minutes. Subjects received a show-up fee of €4.00. In case they were randomly selected to receive additional payments based on our main treatments on memory and representativeness, subjects received €0.50 for each correct answer to the questions on the 50 images (see the full list of questions below). A distraction task on emotional expressions was used in the wave conducted at the University of Cologne and took 90 seconds on average. A distraction task on raven matrices was used in the wave conducted at Bocconi University and took 170 seconds on average. We find no treatment differences between the two waves and hence present in the main text the results of the treatments with distraction by pooling both waves. In the following we also show the non-pooled results.

Images

Screenshots for each type of image used in Study 1 are displayed in Figure B.1. Figures B.2 and B.3 list the target images (orange and blue numbers) and decoy images (blue words or gray shapes, depending on treatment).

Questions

- Q1: The computer randomly chose 1 image from all images that were just shown to you. The chosen image showed a number. What is the likely color of the chosen image? Blue or Orange.
- Q2: The computer randomly chose 1 image from all images that were just shown to you. The chosen image showed a number. What is the probability that this number is orange?

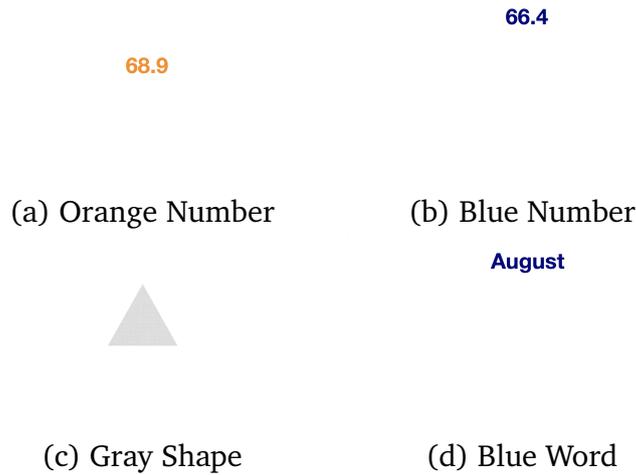


Figure B.1: Examples of image screenshots

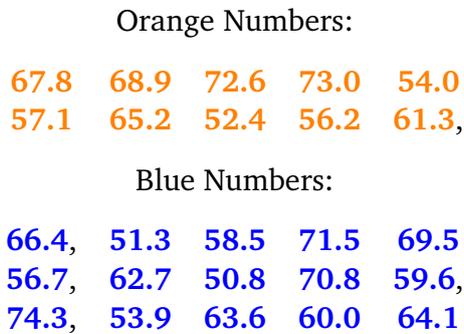


Figure B.2: Full set of target images shown to participants

- Q3: How many orange numbers were shown to you?
- Q4: How many blue numbers were shown to you?
- *Blue* treatment only:

AddQ1: How many blue words were shown to you?

AddQ2: How many orange numbers were shown to you?

- *Gray* treatment only:

AddQ1: How many gray shapes were shown to you?

- Q5: The computer randomly chose 1 image from all images that were just shown to you. The chosen image showed a number. What is the probability that this number is blue?

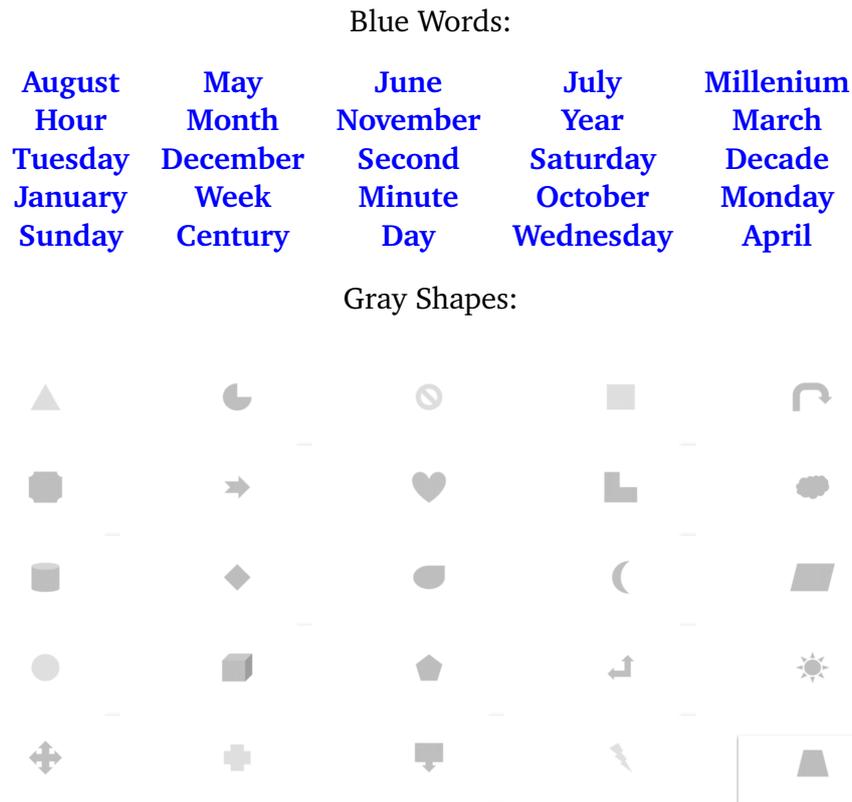


Figure B.3: Decoy images shown to participants depending on treatment

- The question was used only in the second wave of Study 1. Because of a computer error, we have responses to this question only for roughly half of the 2nd wave’s sample.

Distraction tasks After participants were exposed to the 50 images and before they responded to the questions, we employed one of two distraction tasks to wash out participants working memory, the emotion expression task (see Figure B.4 for an example) and the raven matrices task (see Figure B.5 for an example). The goal of the distraction task is to fully engage participants’ working memory in a task orthogonal to the key task at hand (recall of the images). By presenting this distraction task, we likely impose a heavy extraneous cognitive load (see Paas, Renkl, and Sweller 2003), swamping participants’ working memories with new information. If our results persist after the distraction, this suggests a role for recall rather than working memory.

Further results

Our main results of Study 1 reported in Table 3 also hold when (i) providing different statistical tests (see Table B.1) and (ii) looking at our two waves separately (see Table



Figure B.4: One of the tasks of the emotion recognition questionnaire used to distract participants in wave 1 of Study 1.

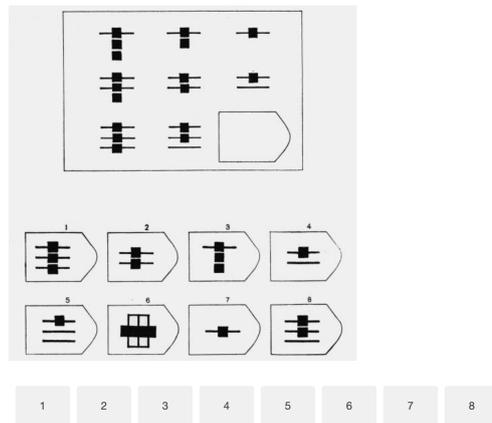


Figure B.5: One of the tasks of the raven matrices used to distract participants in wave 2 of Study 1.

B.2 for wave 1 and Table B.3 for wave 2).

Table B.1: Regression estimates of treatment effects in Study 1

	Logit: Y=1 if "orange is more likely"	Logit: Y= 1 if more orange numbers recalled	OLS Y= Orange numbers recalled	OLS Y= Blue numbers recalled	OLS Y= Ratio of orange to total numbers recalled	OLS Y= Probability that a randomly-drawn number is orange	OLS Y= Probability that a randomly-drawn number is blue
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1 if <i>blue</i>	.8970*** (.1498)	.7741*** (.1564)	.3807 (.6011)	- 1.200* (.6828)	.0412*** (.0116)	.0613*** (.0134)	-.0296 (.0332)
Wave dummy	yes	yes	yes	yes	yes	yes	no
Constant	-.8008*** (.1270)	-1.070*** (.1341)	13.57*** (.4990)	16.29*** (.5668)	.4575*** (.0096)	.4466*** (.0111)	.5498*** (.0253)
Observations	790	790	790	790	790	790	117
Adj./Ps. R ²	0.04	0.03	0.00	0.00	0.01	0.03	0.00

Table B.2: Regression estimates of treatment effects in Study 1b for wave 1

	OLS: Y=1 if “orange is more likely”	OLS: Y= 1 if more orange numbers recalled	0.5-Q-Reg Y= Orange numbers recalled	0.5-Q-Reg Y= Blue numbers recalled	0.5-Q-Reg: Y= Ratio of orange to total numbers recalled	0.5-Q-Reg: Y= Probability that a randomly-drawn number is orange
	(1)	(2)	(3)	(4)	(5)	(6)
1 if <i>blue</i>	.2445*** (.0465)	.1615*** (.0452)	0 (.9386)	0 (.7426)	.0455** (.0217)	.05* (.0277)
Constant	.3049*** (.0322)	.2601*** (.0313)	12*** (.6488)	15*** (.5132)	.4545*** (.0150)	.45*** (.0192)
Observations	427	427	427	427	427	427
(Ps.) R^2	0.05	0.03	0.00	0.00	0.01	0.01

Table B.3: Regression estimates of treatment effects in Study 1b for wave 2

	OLS: Y=1 if “orange is more likely”	OLS: Y= 1 if more orange numbers recalled	0.5-Q-Reg Y= Orange numbers recalled	0.5-Q-Reg Y= Blue numbers recalled	0.5-Q-Reg: Y= Ratio of orange to total numbers recalled	0.5-Q-Reg: Y= Probability that a randomly-drawn number is orange	0.5-Q-Reg: Y= Probability that a randomly-drawn number is blue
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1 if <i>blue</i>	.1954*** (.0504)	.1728*** (.0480)	0 (.6960)	− 2** (.7859)	.0806*** (.0192)	.1*** (.0251)	−.1*** (.0373)
Constant	.2840*** (.0368)	.2189*** (.0351)	10*** (.5088)	14*** (.5745)	.4194*** (.0141)	.4*** (.0184)	.6*** (.0285)
Observations	363	363	363	363	363	363	117
(Ps.) R^2	0.04	0.03	0.00	0.01	0.02	0.03	0.02

B.1.1 Study 1b

We also ran Study 1 in the lab and on MTurk without distraction between the image presentation and questions and without eliciting Q2, called Study 1b. Here is a link to an online version of the two treatments of Study 1b:

https://unikoelnwiso.eu.qualtrics.com/jfe/form/SV_6PaEfC4u67hWp1P

Procedural details and data collection

We conducted Study 1b in three waves. Wave 1 was conducted in February of 2018 with MTurk and a sample of 337 participants. Our *blue* and *gray* treatments were accompanied by an unrelated intertemporal choice. The entire experiment lasted for around 10 minutes. Participants received a \$1.00 show-up fee. A computer-based coin toss determined randomly whether subjects would receive additional payments based on the *blue* and *gray* treatments or on the unrelated intertemporal choice. In case subjects received additional payments based on the former, one of all participants was randomly chosen to receive \$20.00 for each correct answer to Q1, Q3, Q4 and AddQ1-2, while all remaining participants received \$0.20 for each correct answer.

We then successfully replicated our findings in two more waves, one in the laboratory of the University of Cologne in March of 2018 (N=483) and another with MTurk in March 2018 (N=193). In our first replication, we tested whether our results are robust to moving from MTurk to the laboratory. The laboratory offered us more control on the image display, because we could ensure equally stable internet connection and computing power for each participant. Like in the first wave, the *blue* and *gray* treatments were accompanied by an unrelated intertemporal choice. The entire lab experiment also took 10 minutes. Subjects received a show-up fee of € 4.00. One participant per experimental session (consisting of 26 to 32 participants) was randomly selected to receive additional payments based on the intertemporal choice task. All remaining participants received € 0.50 for each correct answer to Q1, Q3, Q4 and AddQ1-2 of the *blue* and *gray* treatments.

In our final replication, we tested whether our results are robust to conducting the *blue* and *gray* treatments without being accompanied by unrelated intertemporal choice tests. The experiment of the third wave took below 7 minutes. Subjects received a \$0.50 show-up fee and \$0.20 for each correct answer to Q1, Q3, Q4, and AddQ1-2.

Questions

In order of exposure to participants:

- Q1: The computer randomly chose 1 image from all images that were just shown to you. The chosen image showed a number. What is the likely color of the chosen image? Blue or Orange.
- Q3: How many orange numbers were shown to you?
- Q4: How many blue numbers were shown to you?
- *Blue* treatment only:

AddQ1: How many blue words were shown to you?

AddQ2: How many orange numbers were shown to you?

- *Gray* treatment only:

AddQ1: How many gray shapes were shown to you?

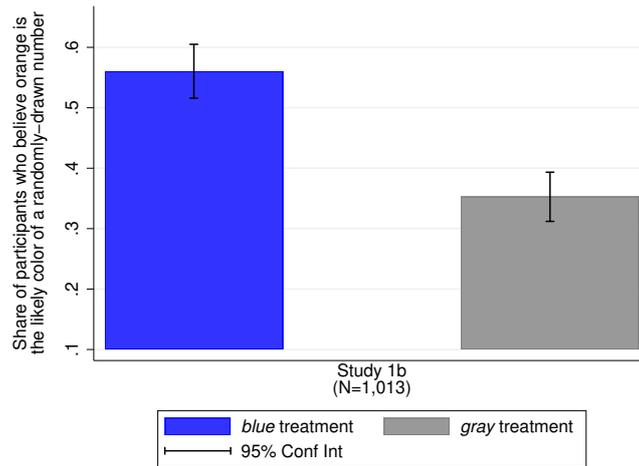


Figure B.6: Share of participants who believe that the likely color of a randomly-drawn number is orange for the *blue* and *gray* treatments of Study 1.

Table B.4: Regression estimates of treatment effects in Study 1b

	OLS: Y=1 if “orange is likely”	OLS: Y= 1 if more orange numbers recalled	0.5-Q-Reg Y= Orange numbers recalled	0.5-Q-Reg Y= Blue numbers recalled	0.5-Q-Reg: Y= Share of orange to total numbers recalled
	(1)	(2)	(3)	(4)	(5)
1 if <i>blue</i>	.2060*** (.0307)	.1379*** (.0302)	0 (.4187)	- 2*** (.6427)	.0556*** (.0124)
MTurk dummy	yes	yes	yes	yes	yes
Wave dummies	yes	yes	yes	yes	yes
Constant	.3305*** (.0302)	.3043*** (.0259)	10*** (.3598)	14*** (.5524)	.4444*** (.0107)
Observations	1,013	1,013	1,013	1,013	1,013
Adj./Ps. R^2	0.04	0.02	0.01	0.01	0.02

Main Results of Study 1b

Further results

The results are summarized in Figure B.6 and Table B.4 and show strong support for Prediction 1. Column (1) in Table B.4 reports an OLS regression of a response dummy (1 if “orange is likely”) on a treatment dummy (1 if *blue*), that amounts to comparing the average share of participants who said orange is likely in the *gray* treatment versus the *blue* treatment. As shown in Figure B.6, that share increases 20.8pp from the *gray* treatment to the *blue* treatment (35.3% to 56%, significant at 1% level).

Column (2) reports the results of an OLS regression of a response dummy that takes

value 1 if the participant reported more orange numbers in Q3 than blue numbers in Q4, which is an alternative measure of each participants' belief about the likely color. Consistent with Column (1), there is a 13.8pp increase in the share of participants who recalled more orange than blue numbers in the *blue* treatment (30.4% to 44.2%, significant at 1% level). Across Columns (1) and (2), the treatment dummy coefficients are close. In fact, answers in Q1, Q3 and Q4 are consistent for roughly 90% of participants. Columns (3) and (4) show how the median quantity of recalled orange and blue numbers depends on the treatment. Responses in the *gray* treatment are quite accurate, as indicated by the constant term. In the *blue* treatment, participants retrieve fewer blue numbers, consistent with interference from blue words.

Finally, based on answers to Q3 and Q4, we can compute the ratio of orange numbers to total numbers recalled. Column (5) shows that, as predicted by the model, participants recalled on average a significantly higher share of orange numbers in the *blue* treatment (50% versus 44.5%). Given that the recalled share of orange numbers is high at baseline, this average increase can have a large effect on the number of participants who say orange is the likely color.

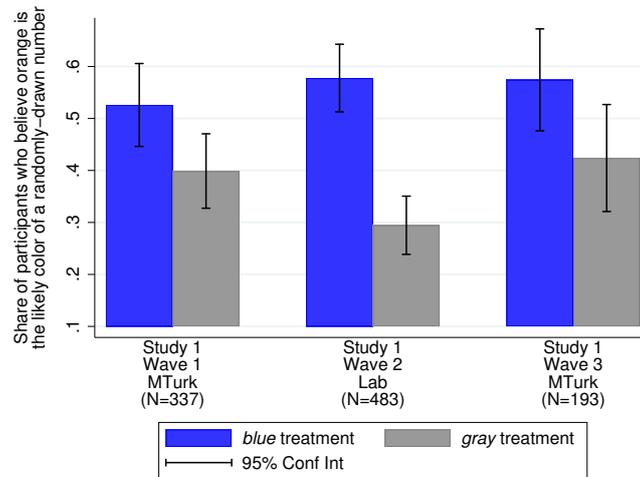


Figure B.7: Share of participants who believe that the likely color of a randomly-drawn number is orange for the *blue* and *gray* treatments for each wave of Study 1b.

Result per wave Figure B.6 and Table B.4 presented the main findings of Study 1b on the outcomes Q1, Q3, and Q4 which provided support for the predictions of our model.

Figure B.7 shows that the share of participants who recalled more orange than blue numbers is larger in the *blue* treatment than in the *gray* treatment for each wave. These differences of 12.71pp, 28.32pp, and 15.03pp for waves 1, 2, and 3, respectively, are significant in OLS regressions (*p*-values of 0.020, <0.001, and 0.037, respectively). Note

that the treatment effect for wave 2—which was conducted in the laboratory—is significantly larger than the pooled treatment effect of waves 1 and 3—which were conducted with MTurk—in a OLS difference-in-differences regression (p -value of 0.018 for the difference-in-differences estimate). Thus, in the laboratory we find a greater treatment effect

Tables B.5, B.6, and B.7 show that the results of Table B.4 also hold when focusing only on wave 1, wave 2, or wave 3, respectively. Only the treatment effects on the median orange numbers recalled and median blue numbers recalled do not replicate for wave 3, which has the smallest sample of the three waves.

Table B.5: Regression estimates of treatment effects in Study 1b for wave 1

	OLS: Y=1 if “orange is likely”	OLS: Y= 1 if more orange numbers recalled	0.5-Q-Reg Y= Orange numbers recalled	0.5-Q-Reg Y= Blue numbers recalled	0.5-Q-Reg Y= Share of orange to total numbers recalled
	(1)	(2)	(3)	(4)	(5)
1 if <i>blue</i>	.1271** (.0542)	.0758 (.0527)	0 (.3916)	− 2** (.9791)	.0556*** (.0100)
Constant	.3989*** (.0366)	.3333*** (.0356)	10*** (.2647)	14*** (.6618)	.4444*** (.0144)
Observations	337	337	337	337	337
Adj./Ps. R^2	0.02	0.01	0.00	0.01	0.02

Table B.6: Regression estimates of treatment effects in Study 1b for wave 2

	OLS: Y=1 if “orange is likely”	OLS: Y= 1 if more orange numbers recalled	0.5-Q-Reg Y= Orange numbers recalled	0.5-Q-Reg Y= Blue numbers recalled	0.5-Q-Reg Y= Share of orange to total numbers recalled
	(1)	(2)	(3)	(4)	(5)
1 if <i>blue</i>	.2832*** (.0433)	.1754 (.0434)	0 (.4541)	− 2*** (.7241)	.0556*** (.0175)
Constant	.2946*** (.0366)	.2868*** (.0296)	10*** (.3099)	14*** (.4942)	.4444*** (.0119)
Observations	483	483	483	483	483
Adj./Ps. R^2	0.08	0.01	0.00	0.02	0.02

Table B.7: Regression estimates of treatment effects in Study 1b for wave 3

	OLS: Y=1 if “orange is likely”	OLS: Y= 1 if more orange numbers recalled	0.5-Q-Reg Y= Orange numbers recalled	0.5-Q-Reg Y= Blue numbers recalled	0.5-Q-Reg Y= Share of orange to total numbers recalled
	(1)	(2)	(3)	(4)	(5)
1 if <i>blue</i>	.1503** (.0716)	.1521** (.0692)	2 (1.612)	3** (1.408)	.0556** (.0100)
Constant	.4239*** (.0518)	.2935*** (.0501)	10*** (1.166)	12*** (1.019)	.4444*** (.0190)
Observations	193	193	193	193	193
Adj./Ps. R^2	0.02	0.02	0.00	0.01	0.03

Different tests We show in Table B.8 that our results presented in Table B.4 are robust to using different statistical tests (Logit regressions instead of OLS regressions for outcome measures of Columns 1 and 2 as well as OLS regressions instead of 0.5 quantile regressions for Columns 3, 4, and 5). Participants are significantly more likely to believe that a randomly-drawn image is likely to be orange (Column 1). Participants are significantly more likely to recall more orange than blue numbers in the *blue* treatment than in the *gray* treatment (Column 2). Additionally, participants state a significantly greater average share of orange numbers recalled to total amount of images recalled in the *blue* treatment than in the *gray* treatment (Column 5). Columns 3 and 4 show that subjects recall on average more orange numbers and less blue numbers in the the *blue* treatment than in the *gray* treatment, however these differences are not significant.

Table B.8: Robustness of regression estimates of treatment effects in Study 1

	Logit: Y=1 if “orange is likely”	Logit: Y= 1 if more orange numbers recalled	OLS Y= Orange numbers recalled	OLS Y= Blue numbers recalled	OLS Y= Share of orange to total numbers recalled
	(1)	(2)	(3)	(4)	(5)
1 if <i>blue</i>	.8450*** (.0307)	.5948*** (.0302)	.4667 (.4130)	− .1399 (.5190)	.0341*** (.0100)
MTurk dummy	yes	yes	yes	yes	yes
Wave dummies	yes	yes	yes	yes	yes
Constant	−.7022*** (.1137)	−.8270*** (.1162)	11.56*** (.3549)	14.06*** (.4461)	.4517*** (.0144)
Observations	1,013	1,013	1,013	1,013	1,013
Adj./Ps. R^2	0.03	0.02	0.01	0.00	0.01

B.1.2 Differences between Study 1 and Study 1b

Table B.9 shows OLS difference-in-difference regressions that test whether the treatment effects of Study 1 (with distraction) and Study 1b (without distraction) differ significantly from each other. The highlighted row shows the difference-in-differences estimates, which are zero or close to zero for all dependent variables in size and do not differ from zero significantly for any of the dependent variables.

Table B.9: Comparing regression estimates of treatment effects between Studies 1 & 1b

	OLS: Y=1 if “orange is more likely”	OLS: Y= 1 if more orange numbers recalled	0.5-Q-Reg Y= Orange numbers recalled	0.5-Q-Reg Y= Blue numbers recalled	0.5-Q-Reg: Y= Share of orange to total numbers recalled
	(1)	(2)	(3)	(4)	(5)
1 if <i>blue</i>	.2060*** (.0307)	.1379*** (.0302)	0 (.4187)	− 2*** (.6427)	.0556*** (.0124)
1 if distraction	− .0549 (.0409)	− .0821** (.0398)	0 (.6253)	0 (.8114)	−.0159 (.0164)
1 if <i>blue</i> & distraction	.0051 (.0460)	.0288 (.0448)	0 (.7041)	0 (.9134)	0 (.0185)
MTurk dummy	yes	yes	yes	yes	yes
Wave dummies	yes	yes	yes	yes	yes
Constant	.3305*** (.0302)	.3043*** (.0259)	10*** (.3598)	14*** (.5524)	.4444*** (.0107)
Observations	1,803	1,803	1,803	1,803	1,803
Adj./Ps. R^2	0.05	0.02	0.01	0.01	0.02

B.2 Study 2

Orange words

Study 2 implements the same design as Study 1 *blue words* treatment, with the one difference that some of the words in the decoy group were shown in orange and the remaining were shown in blue as before.

Procedural details and data collection

We conducted three waves of Study 2’s treatments. The first two waves were conducted in March and May of 2018 with MTurk, $N = 307$ and $N = 1,431$, respectively. In both of these waves, the experiment only consisted of the blue treatment variation with orange words. The experiment lasted for 7 minutes for both waves. Participants received a \$1.00 show-up fee as well as \$0.20 for each correct answer. In the first wave we conducted the treatments with $k = 0, 1, 3, 6$. In the second wave we replicated the first wave and included in addition treatments with $k = 10, 22$.

The third wave was conducted in May of 2018 in the laboratory of Bocconi University with $k = 1, 6$. These treatments were accompanied by an unrelated intertemporal choice—like in the laboratory experiments of Studies 1 and 1b. The entire lab experiment took 10 minutes. Subjects received a show-up fee of €4.00. In case they were randomly selected to receive additional payments based on the *blue* treatments with

$k = 1, 6$, subjects received € 0.50 for each correct answer to the 50 images. We are using the *blue* treatment of the lab experiment of Study 1 and Study 1b as a comparison standard.

Further results

Table B.10 shows our findings when looking at MTurk participants only. Like in Table 4 from the main text that shows our results when looking at lab participants only, we find evidence for Prediction 3. The share of participants who say “orange is more likely” drops by 1pp per orange word added. Turning to data on recall, as orange words are added, participants recall more blue numbers and are less likely to recall seeing more orange numbers.

Table B.10: Regression estimates of treatment effects in Study 2 (MTurk only)

	OLS: Y=1 if “orange is likely”	0.5-Q-Reg: Y= Orange numbers recalled	0.5-Q-Reg: Y= Blue numbers recalled	OLS: Y= 1 if more orange numbers recalled	0.5-Q-Reg: Y= Share of orange to total num- bers recalled
	(1)	(2)	(3)	(4)	(5)
k (number of orange words)	−.0093*** (.0015)	0 (.0488)	.0625* (.0334)	−.0030** (.0015)	−.0005 (.0006)
Constant	.5265*** (.0015)	12*** (.5031)	13.625*** (.3421)	.4116*** (.0156)	.5005*** (.0058)
Observations	1,839	1,839	1,839	1,839	1,839
Adj./Ps. R^2	0.02	0.00	0.00	0.02	0.00

Notes: This table presents estimates of the treatment effect on several outcome measures relating to our model predictions in Study 2. We report on our results from MTurk here. Columns (1) and (5) present OLS regression of response dummies (1 if “orange is likely” in Column (1) and 1 if “recalled more orange than blue numbers” in Column (3)) on k , the number of orange words. Columns (2), (3), (4) and (6) present 0.5-quantile regressions of subjects’ stated probability that a random number is orange, number of recalled orange numbers, number of recalled blue numbers and share of recalled oranges numbers, respectively, on k . * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

B.3 Study 3

Questions

The following questions were asked to participants, in the order presented below. The labelling used here mirrors that in the main text.

Color Cue Treatment

- Screen 1
 - Q1: The computer randomly chose 1 image from all images that were just shown to you. The chosen image showed a number. What is the likely color of the chosen image? Blue or Orange.
- Screen 2
 - Q2: The computer randomly chose 1 image from all images that were just shown to you. The chosen image showed a number. What is the probability that this number is orange?
- Screen 3
 - Q3a: How many blue numbers in small font size were shown to you?
 - Q3b: How many blue numbers in large font size were shown to you?
 - Q4a: How many orange numbers in small font size were shown to you?
 - Q4b: How many orange numbers in large font size were shown to you?

Size Cue Treatment (per screen in order of display)

- Screen 1
 - Q1: The computer randomly chose 1 image from all images that were just shown to you. The chosen image showed a number. What is the likely font size of the chosen image? Small or Large.
- Screen 2
 - Q4: The computer randomly chose 1 image from all images that were just shown to you. The chosen image showed a number. What is the probability that this number is large?
- Screen 3

- Q2a: How many blue numbers in small font size were shown to you?
- Q2b: How many blue numbers in large font size were shown to you?
- Q3a: How many orange numbers in small font size were shown to you?
- Q3b: How many orange numbers in large font size were shown to you?

Procedural details and data collection

We conducted two waves of Study 3. Wave 1 was conducted in May of 2018 in the laboratory of the Bocconi University ($N = 326$). Wave 2 replicated Wave 1 in October of 2018 in the laboratory of the University of Cologne ($N = 321$). The memory treatments were accompanied by an unrelated intertemporal choice—like in the laboratory experiment of Studies 1 and 1b. The entire lab experiment took 10 minutes. Participants received a show-up fee of €4.00. In case they were randomly selected to receive additional payments based on the treatments of Study 3, participants received €0.50 for each correct answer to the questions Q1-Q4.

Further results

We find evidence for our prediction for both Q1 and Q4 in both waves of Study 3, see Figures B.8 and B.9.

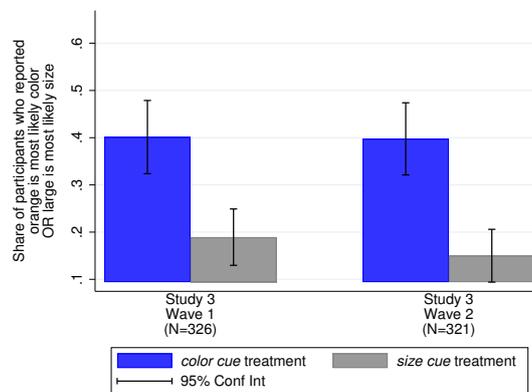


Figure B.8: Share of participants who believe that the color OR size of a randomly-drawn number is most likely orange OR large for the treatments of Study 3.

B.3.1 Within-subject comparison of outcome measures across Studies 1, 2, 3

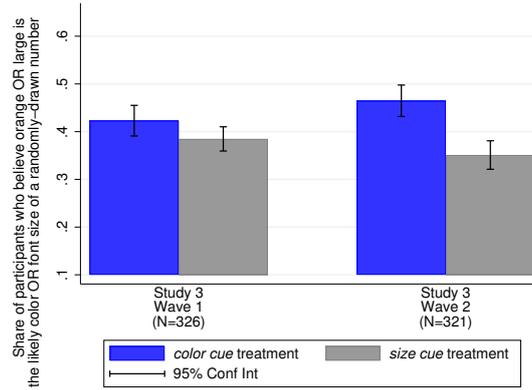


Figure B.9: Participants' belief that a random number is orange OR large for the treatments of Study 3.

Table B.11: Relationship between the share of recalled orange over total numbers and the direct probability measure

	OLS: Y=		
	Probability that a randomly-drawn number is orange		
	Study 1 (1)	Study 2 (2)	Study 3 (3)
Share of recalled orange over total numbers	.8607*** (.0278)	.7284*** (.0628)	.3751*** (.0432)
Wave dummy	yes		yes
Constant	.0611*** (.0139)	.1272*** (.0307)	.2410*** (.0204)
Observations	790	254	644
Adj./Ps. R^2	0.55	0.35	0.10

Notes: This table presents estimates of the relationship between the share of recalled orange over total numbers and the direct probability measure. Columns (1), (2) and (3) present OLS regression of participants' imputed stated probability that a random number is orange on their stated share of recalled orange over total numbers, respectively, for each of study 1, study 2 and study 3. Standard errors are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

C Further Experiments

In the following, we discuss 4 further studies that we conducted. The first study, which we call study 1.A in the main text, builds on study 1 but varies the number of decoy images presented. This study is used in the assessment of the performance of the calibrated model in Section 5. The second study instead varies the composition of the target group. The third and fourth studies vary other design features.

C.1 Study 1.A

Design

In Study 1.A, we vary the number of decoy images, as well as their content (blue words vs gray shapes), in a 4 x 2 design presented in Table C.1.

Table C.1: Study 1.A's *Blue j* Treatment and *Gray j* Treatment

Blue <i>j</i> Treatments (with $j \in \{5, 50, 75, 125\}$)			
	Orange	Blue	Gray
Target data	10 Orange Numbers	15 Blue Numbers	
Decoy data		<i>j</i> Blue Words	

Gray <i>j</i> Treatments (with $j \in \{5, 50, 75, 125\}$)			
	Orange	Blue	Gray
Target data	10 Orange Numbers	15 Blue Numbers	
Decoy data			<i>j</i> Gray Shapes

Questions The following questions were asked to participants, in the order presented below.

- Q1: The computer randomly chose 1 image from all images that were just shown to you. The chosen image showed a number. What is the likely color of the chosen image? Blue or Orange.
- Q2: How many orange numbers were shown to you?
- Q3: How many blue numbers were shown to you?
- *Blue* treatments only:

AddQ1: How many blue words were shown to you?

AddQ2: How many orange words were shown to you?

- *Gray* treatments only:

AddQ1: How many gray shapes were shown to you?

Procedural details and data collection

We conducted two waves of Study 1.A. The first wave was conducted in March of 2018 with MTurk and a sample of 800 participants featuring all treatment cells of the 4×2 design. We then replicated the *blue* treatments with $j = 5, 50,$ and 75 decoys in May of 2018 with MTurk and a sample of 592 participants. In both waves, the experiment consisted only of our memory treatments. The experiment lasted for 7 minutes. Participants received a \$1.00 show-up fee as well as \$0.20 for each correct answer to the questions on the 50 images.

Results

Figure C.1 shows that we find no treatment effect when the blue and orange numbers are displayed along 5 decoys (in the *blue* and the *gray* treatments). For 50, 75, and 125 decoys, however, we do find significant treatment effects that resemble our findings of Study 1.

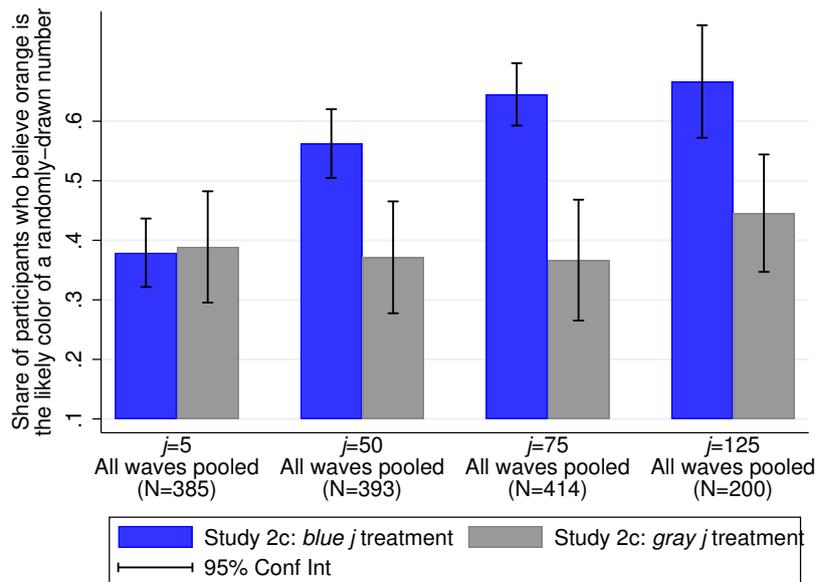


Figure C.1: Share of participants who state orange is the likely color, by treatment.

C.2 Study 2.A

Design

In Study 2.A, we vary likelihood ratios by changing the composition of images within our target data, as well as their content (blue words vs gray shapes), in a 3×2 design presented in Table C.2.

Table C.2: Study 2.A's *Blue i* Treatment and *Gray i* Treatment

Blue <i>i</i> Treatments (with $i \in \{13, 20, 25\}$)			
	Orange	Blue	Gray
Target data	25- <i>i</i> Orange Numbers	<i>i</i> Blue Numbers	
Decoy data		25 Blue Words	

Gray <i>i</i> Treatments (with $i \in \{13, 20, 25\}$)			
	Orange	Blue	Gray
Target data	25- <i>i</i> Orange Numbers	<i>i</i> Blue Numbers	
Decoy data			25 Gray Shapes

Questions The following questions were asked to participants, in the order presented below.

- Q1: The computer randomly chose 1 image from all images that were just shown to you. The chosen image showed a number. What is the likely color of the chosen image? Blue or Orange.
- Q2: How many orange numbers were shown to you?
- Q3: How many blue numbers were shown to you?
- *Blue* treatments only:

AddQ1: How many blue words were shown to you?

AddQ2: How many orange words were shown to you?

- *Gray* treatments only:

AddQ1: How many gray shapes were shown to you?

Predictions

Our prediction is that as we increase the share of blue numbers, the likelihood that participants recall a randomly-chosen number as orange should decrease in the *blue* treatment. Thus, we expect a smaller treatment effect (difference between *blue* and *gray*) as

the share of blue numbers increases. Denote the number of blue numbers as i , so that the number of orange numbers is $25-i$ and assume $c > 0$. Then, representativeness-based recall yields the following prediction:

As the share of blue to orange numbers increases, the share of participants who believe that the likely color of a random number is orange should decrease in the *blue* treatment, because the assessed probability that a random number is blue increases, formally $\tilde{P}(b|n)_{blue\ i} \geq \tilde{P}(b|n)_{blue\ i'}$ for $i > i'$.

Procedural details and data collection

We conducted two waves of Study 2.A. The first wave was conducted in March of 2018 via MTurk with 601 participants. The experiment consisted only of our memory treatments. The experiment lasted for 7 minutes. Participants received a \$1.00 show-up fee as well as \$0.20 for each correct answer. We then replicated these results in the laboratory of the University of Cologne with 516 participants. The memory treatments were accompanied by an unrelated intertemporal choice—like in the laboratory experiment of Study 1. The entire lab experiment took 10 minutes. Participants received a show-up fee of €4.00. In case they were randomly selected to receive additional payments based on our treatments on memory and representativeness, participants received €0.50 for each correct answer to the questions on the 50 images.

Results

Our findings provide evidence for our prediction. When the share of blue to orange numbers is increased, participants' *blue* to *gray* treatment difference in their assessed probability that a random number is orange decreases. The treatment effect on the share of participants believing that “orange is more likely” is 15.3pp for 13 blue numbers and 12 orange numbers. The treatment effect reduces to 10pp for 20 blue and 5 orange numbers and to 6pp for 23 blue and 2 orange numbers. While all three treatment effects are (at least weakly) significantly different from zero, the former one is larger than the latter two. OLS regressions show that this difference in treatment effects is at most weakly significantly different from zero when comparing the 13 blue numbers and 12 orange numbers case with the 23 blue numbers and 2 orange numbers case as well as when comparing the 13 blue numbers and 12 orange numbers case with the pooled cases of 23 blue numbers and 2 orange numbers as well as 20 blue numbers and 5 orange numbers. Column (1) of Table C.3 shows the results of an OLS regression of the latter result. The weakly significant interaction term (Row (3)) implies the discussed difference in treatment effects. Column (2) of Table C.3 shows that the difference in treatment effects increases in size and significance when the main treatments of Study 1 are included:

The treatment effect for treatments with 20 or 23 blue numbers is significantly smaller than for treatments with 13 or 15 blue numbers.

Table C.3: Regression estimates of treatment effects in Study 2.A

	Including only Study 2.A OLS: Y=1 if “Orange is more likely” (1)	Including Study 1 & 2.A OLS: Y=1 if “Orange is more likely” (2)
1 if <i>blue</i>	.1669*** (.0411)	.1957*** (.0237)
1 if <i>i</i> (Blue Numbers) ≥ 20	-.3044*** (.0354)	-.2897*** (.0397)
1 if <i>blue</i> & <i>i</i> ≥ 20	-.0904* (.0499)	-.1192*** (.0656)
MTurk dummy	yes	yes
wave dummies	-	yes
Constant	.3302*** (.0325)	.3153*** (.0300)
Observations	1,117	2,130
Adj. R^2	0.18	0.15

Notes: For Study 2.A, wave dummy and MTurk dummy are collinear

C.3 Fontsize Study

The design of our fontsize Study is similar to that of Study 1, except that images vary in font size as opposed to color.

Design

The two between-subjects treatments of our fontsize Study follow the same structure as our baseline treatments of Study 1:

First, participants are told they will see questions that are incentivized for accuracy about a sequence of 50 abstract images displayed to them during the experiment.

Second, participants see the 50 images, each of which appears on separate screens for short moments of time and in random order. The 50 images vary along two features. As in previous Studies, images can be a number or a word. We continue to refer to numbers as targets and words as decoys. The second feature is the font size of the object, which

Table C.4: Fontsize Study's *small* treatment and *large* treatment

<i>Small (Font Size) Treatment</i>		
	Large	Small
Target data	10 Large Numbers	15 Small Numbers
Decoy data		25 Small Words

<i>Large (Font Size) Treatment</i>		
	Large	Small
Target data	10 Large Numbers	15 Small Numbers
Decoy data	25 Large Numbers	

can be large or small. Table C.4 shows which types of images participants were shown for each of the two treatments. Example screenshots for each kind of image are displayed in Figure C.2.

Third, participants face questions on the targets which require them to recall the observed sequence of images as well as on the decoys images. The questions – presented here in the same order as shown to participants in the experiment—were:

- Q1: The computer randomly chose 1 image from all images that were just shown to you. The chosen image showed a number. What is the likely size of the chosen image? Large or Small.
- Q2: How many large numbers were shown to you?
- Q3: How many small numbers were shown to you?
- Q4: How many small words were shown to you?
- Q5: How many large words were shown to you?

Predictions

Our prediction is that as we change the decoy images from being large words to small words, the degree of participants who believe that a randomly-chosen number was shown in large font size should be greater. Because large numbers are representative only in the *small* treatment, interference-based recall inhibits recall of small numbers more in the *small* treatment than in the *large* treatment. Thus, we expect more participants to state that the likely font size is large in the *small* treatment than in the *large* treatment.

53.7

(a) Large Number

(b) Small Number

Week

Decade

(c) Large Word

(d) Small Word

Figure C.2: Examples of images shown to participants

Procedural details and data collection

We conducted the fontsize Study in December of 2017 with MTurk and a sample of 374 participants. Our *large* and *small* treatments were accompanied by unrelated intertemporal choices. The entire experiment lasted for around 13 minutes. Participants received a \$1.00 show-up fee. A computer-based coin toss determined randomly whether subjects would receive additional payments based on the *blue* and *gray* treatments or on the unrelated intertemporal choice. In case subjects received additional payments based on the former, one of every 100 participants was randomly chosen to receive \$20.00 for each correct answer to Q1-Q5, while all remaining participants received \$1 for each correct answer.

Results

We find that 22.4% of participants believe that the likely font size of a randomly-drawn number is large in the *large* treatment. However, in the *small* treatment, 32% of participants believe that the likely font size of a randomly-drawn number is large. This difference is consistent with our model's prediction and is significantly greater than zero in an OLS regression, see Table C.5 Column 1.

Mirroring our results from Study 1, we also find that a greater share of participants recalls more large than small numbers in the *small* treatment than in the *large* treatment, Column 2 of Table C.5. However, we also find that the median amount of recalled large

numbers as well as the median amount of recalled small numbers do not differ across treatments.

Table C.5: Regression estimates of treatment effects in the fontsize Study

	OLS: Y=1 if “orange is likely”	OLS: Y= 1 if more large numbers recalled	0.5-Q-Reg Y= Large numbers recalled	0.5-Q-Reg Y= Small numbers recalled	0.5-Q-Reg Y= Share of large to total numbers recalled
	(1)	(2)	(3)	(4)	(5)
1 if <i>small</i>	.1006** (.0460)	.0883** (.0402)	0 (.8609)	0 (.9589)	.0274* (.0157)
Constant	.2240*** (.0329)	.1421*** (.0287)	8*** (.6152)	15*** (.6852)	.3571*** (.0112)
Observations	374	374	374	374	373
Adj./Ps. R^2	0.01	0.01	0.00	0.00	0.00

C.4 Study ProbNumber

Design

The two between-subjects treatments of our Study ProbNumber follow the same structure as our baseline treatments of Study 1:

First, participants are told they will be asked questions that are incentivized for accuracy about a sequence of 50 abstract images displayed to them during the experiment.

Second, participants see the 50 images that appear on separate screens for short moments of time and in random order. The 50 images vary along two features. The first feature is the object category: each object can be a number, a word or shape. The second feature is color, which can be blue or orange. Table C.6 shows which types of images participants were exposed for each of the two treatments. Example screenshots for each kind of image are displayed in Figure C.3. In the ProbNumber study, blue objects are the target images and the orange objects are the decoys.

Third, participants face one question on the targets which require them to recall the observed sequence of images:

- Q1: The computer randomly chose 1 image from all images that were just shown to you. The chosen image showed a blue object. What is the probability that the chosen image is a number?

Table C.6: Study ProbNumber’s *numbers* treatment and *shapes* treatment

<i>Numbers Treatment</i>		
	Decoy data	Target data
Numbers	10 Orange Numbers	15 Blue Numbers
Non-numbers		25 Blue Words

<i>Shapes Treatment</i>		
	Decoy data	Large data
Numbers		15 Blue Numbers
Non-numbers	10 Orange Shapes	25 Blue Words

Prediction

Our prediction is that as we change the decoy images from being orange shapes to orange numbers, the likelihood that participants recall a randomly-chosen blue object as being a number should be lower. Because orange numbers are representative only in the *numbers* treatment, interference-based recall inhibits recall of blue numbers more in the *numbers* treatment than in the *shapes* treatment. Thus, we expect participants to state lower probabilities that a randomly-chosen blue object is a number in the *numbers* treatment than in the *shapes* treatment.

Procedural details and data collection

We conducted the ProbNumber Study in March of 2018 with MTurk and a sample of 304 participants. The experiment consisted only of the *numbers* and *shapes* treatments. The entire experiment lasted for around 7 minutes. Participants received a \$1.00 show-up fee and \$1.00 for a correct answer to Q1.

Results

Consistent with our prediction, we find that participants believe that a randomly-drawn blue object is a number to a greater extent in the *shapes* treatment—where orange shapes are unlikely to interfere with the recall of blue numbers—than in the *numbers* treatment—where orange numbers are predicted to interfere with the recall of blue numbers.

The median probability that a blue object is a number is 50% in the *shapes* treatment and 40% in the *numbers* treatment. This difference is significant at the 1% level in a

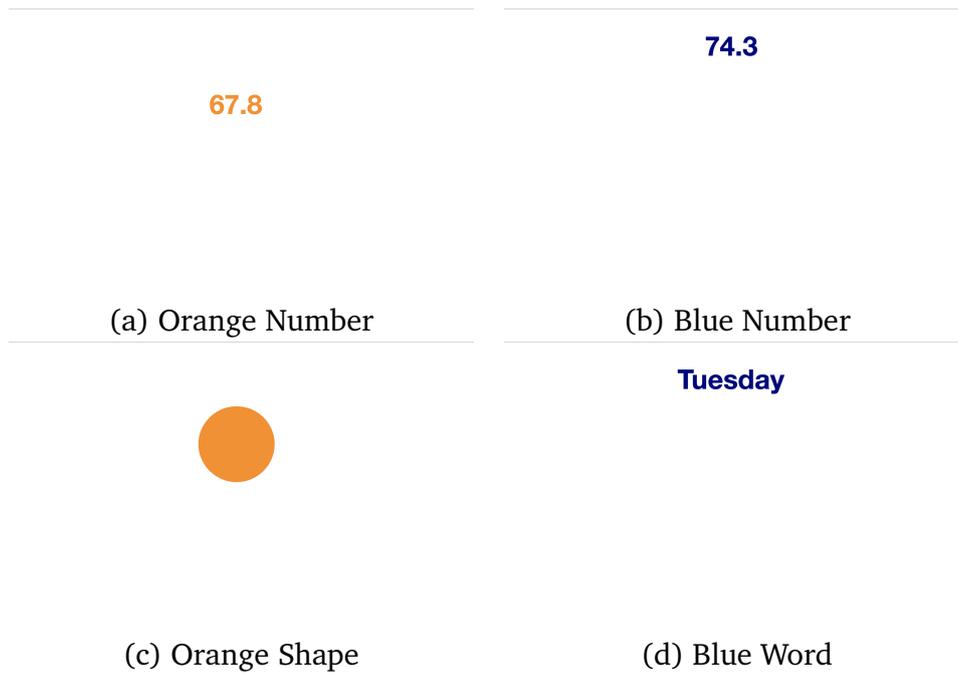


Figure C.3: Examples of images shown to participants in Study ProbNumbers

0.5 quantile regression. We find a similar treatment effect when looking at the mean instead of the median. The mean probability that a blue object is a number is 48% in the *shapes* treatment and 42% in the *numbers* treatment. This difference is significant at the 5% level in an OLS regression.

The results of ProbNumber provide further evidence for how interference-based recall drives distortions of probabilistic judgements.