A Cognitive Foundation for Perceiving Uncertainty[a]

Josh Hascher[b], Alex Imas[c], Michael Ungeheuer[d], and Martin Weber[e]

January 29, 2024

Abstract

We propose a framework where perceptions of uncertainty are driven by the interaction between cognitive constraints and the way that people learn about it—whether information is presented sequentially or simultaneously. People can learn about uncertainty by observing the distribution of outcomes all at once (e.g., seeing a stock return distribution) or sampling outcomes from the relevant distribution sequentially (e.g., experiencing a series of stock returns). Limited attention leads to the overweighting of unlikely but salient events—the dominant force when learning from simultaneous information—whereas imperfect recall leads to the underweighting of such events—the dominant force when learning sequentially. A series of studies show that, when learning from simultaneous information, people are overoptimistic about and are attracted to assets that exhibit large but unlikely outperformance. However, they overwhelmingly select more consistently outperforming assets when learning the same information sequentially, and this is reflected in beliefs. The entire 40-percentage point preference reversal is driven by limited attention and memory; manipulating these factors completely eliminates the effect of the learning environment on choices and beliefs, and can even reverse it.

Keywords: Choice Under Risk, Bounded Rationality, Description-Experience Gap, Attention, Salience, Memory

[a]Support from the Research Council of Finland is gratefully acknowledged. We are grateful to Peter Andre and Aislinn Bohren for valuable comments. Thanks to Haiying Lu and Kaushal Addanki for excellent research assistance.

[b]Josh Hascher: University of Chicago, josh.hascher@chicagobooth.edu.

[c]Alex Imas: University of Chicago, NBER, CESifo, HCEO. alex.imas@chicagobooth.edu.

[d]Michael Ungeheuer: Aalto University, michael.ungeheuer@aalto.fi.

[e]Martin Weber: University of Mannheim and the Centre for Economic Policy Research, London EC1V 3PZ, United Kingdom, weber@bank.BWL.uni-mannheim.de.

1 Introduction

People can learn about the uncertainty involved in their choices through different ways. In some cases, they have access to information about all the potential outcomes and their likelihoods simultaneously. For example, price charts describe the historical returns of a stock and serve as a proxy for the underlying risk distribution; similarly, information about the potential outcomes and their odds are readily available for sports betting, casino gambling, and public lotteries. In many other cases, however, people make risky choices without accessing such information and learn about the relevant distribution by sequentially sampling from it; the riskiness of a new business relationship or a stock can also be judged by observing and experiencing the relevant outcomes. In standard theory, the method through which one learns about uncertainty should not impact decision-making. As long as information is held constant, choices and beliefs about uncertainty should be the same regardless of whether one learns the information simultaneously or sequentially.

However, cognitive constraints may change people's perceptions of uncertainty as a function of the learning environment. We propose a framework that nests well-studied limits on attention and memory and derive predictions on how these limits affect beliefs and choices depending on whether information is presented simultaneously or sequentially.[1] In the case of learning from simultaneous information, memory plays little role and bottom-up channeled attention leads to the overweighting of unlikely but salient events. This can generate a preference for prospects that are most likely to underperform, but when they do outperform, they do so by a lot. A strict preference for such unlikely-outperformers can be interpreted as the overweighting of rare events. When learning through sequential sampling, however, imperfect memory recall implies that an infrequent event will be misremembered. This can shift preferences towards prospects that outperform most frequently—behavior that looks like the underweighting of rare events. Importantly, because the behavioral predictions are derived from basic cognitive processes, the learning environment is expected to change people's perceptions and mental representations of the risky prospects—learning from information presented simultaneously versus sequentially will lead to different *beliefs* about the relevant distributions.

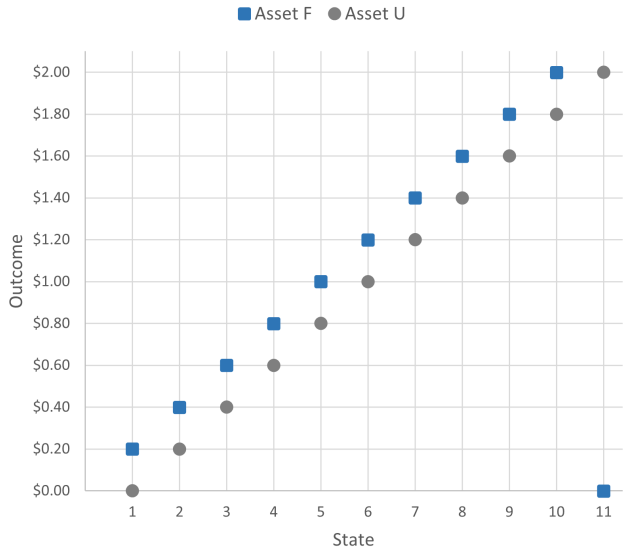
[1] As outlined in greater detail in Section 2, our framework builds on the models in the economics and psychology literature on limited attention, salience, and bounded memory (Bordalo et al., 2012; Li and Camerer, 2022; Bruce and Tsotsos, 2009; Anderson et al., 1998; Bordalo et al., 2022a).

We test the predictions of the framework directly. People faced a choice between two risky assets with the *same* underlying distributions of outcomes. Outcomes were organized such that one asset (Asset F) outperformed in most situations while the other (Asset U) outperformed in one situation by a lot. The distributions of outcomes for Assets U and F are presented in Figure 1 below. Here, a strict preference for Asset U can be interpreted as the overweighting of unlikely states, while a strict preference for Asset F can be interpreted as the underweighting of unlikely states.

Our framework predicts that the salience of the situation with the large performance difference will shift choices towards the unlikely-outperforming Asset U when the distributions are presented simultaneously. On the other hand, the chance of misremembering any given outcome will push people to select Asset F when they learn the same exact information sequentially. Indeed, Asset U was selected 58% of the time in the former environment and Asset F was selected 85% of the time in the latter environment—a more than 40 percentage point gap in choices. These results are summarized in the Baseline conditions reported in Figure 2 below. Data on people's beliefs shows that these choices were a reflection of changes in people's perceptions of the prospects as a function of the learning environment: they had overoptimistic beliefs about the performance of Asset U (Asset F) when learning about the distributions simultaneously (sequentially). These results are particularly surprising because the underlying distributions of the assets were identical; cognitive constraints generated biased beliefs in both learning environments, but the direction of the bias reversed depending on whether information was presented simultaneously or sequentially.

Our proposed mechanism is inherently context-specific and comparative. A state is salient and thus overweighted based on *relative* differences in payoffs and visual factors within the choice set; a state is more or less likely to be misremembered depending on the number of other states and the similarity of elements within the choice set. Models where prospects are evaluated in isolation (e.g., prospect theory) would not predict a difference in choices. To test our framework directly, we exogenously manipulated context-specific factors that influence attention and memory, and studied the impact on choices and beliefs. In the simultaneous information environment, decreasing the salience of the unlikely state where Asset U outperformed shifted choices from 58% of people selecting Asset U ("Baseline") to 66% selecting Asset F ("No Salience Bias"), a preference reversal of 24 percentage points. Importantly, this manipulation was predicted to impact relative salience

Figure 1: Distributions of Outcomes: Frequently-outperforming Asset F versus unlikely-outperforming Asset U. All states are equally likely (1-in-11). The marginal outcome distributions of the two assets are exactly the same (discrete uniform on the $0.20 grid from $0.00 to $2.00), only shifted by one state.



without affecting any underlying properties of the assets, including the state-specific outcomes.[2] The preference reversal against the unlikely-outperforming Asset U thus provides evidence for bottom-up attention as a driver for the overweighting of unlikely events in this and previous work (e.g., Kahneman and Tversky (1979)).

On the other hand, decreasing the scope for imperfect recall in the sequential information environment lowered the preference for Asset F by more than 20 percentage points.[3] Together, these two manipulations completely erased the gap in both choices *and* beliefs, which suggests that neither preference primitives for low-probability events (e.g., probability weighting) nor something inherent to learning from sequential information (e.g., reinforcement learning) is responsible for the gap.

Besides repetition, an important factor governing memory recall is feature-based similarity (Kahana, 2012; Anderson et al., 1998). When making a choice between two risky assets, people are more likely to recall outcomes that share similar features with an option when forming beliefs about the relevant distributions. However, if outcomes share similar features either across states or across assets, this creates the potential for *interference*, where a person attempting to recall the outcome

[2] As discussed further in Section 3, we validated our assumption on shifting salience using Graph Based Saliency tools from cognitive psychology.

[3] We did this in two ways: having people keep a record of the sample ("No Recall Bias"), and reducing the length of the sample to minimize memory requirements ("Minimal Requirement"). Both generated a similar reduction in the selection of Asset F.

of an asset in a given state may incorrectly recall an outcome from a different state or asset if the features are sufficiently similar. To test the interference mechanism directly, we included a third Asset D across two sequential information conditions: one where the third asset shared similar but payoff-irrelevant features with Asset U and another where it shared similar features with Asset F. The third asset's payoffs were similar across conditions and generally better than those of the others. Importantly, people could not choose this third asset in either condition: the choice was always between Asset U and F as in the Baseline study.

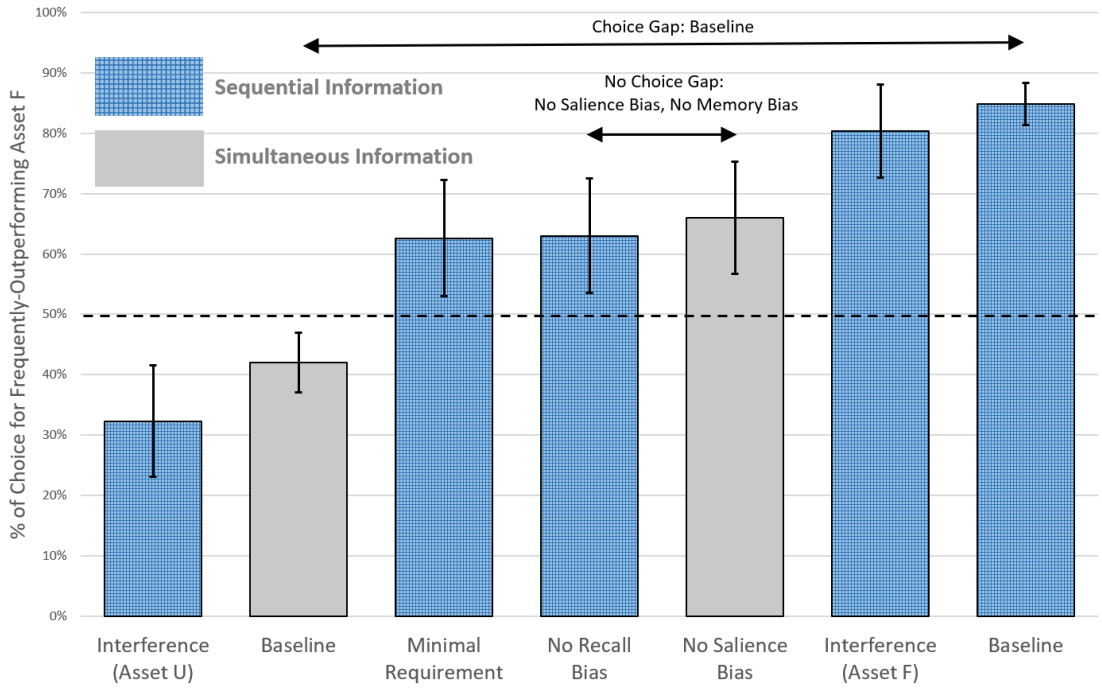
Our framework predicts that despite the choice and learning environment remaining the same, people will be more likely to choose Asset U when the third asset is similar to it—because they incorrectly recall the outcomes of Asset D when forming beliefs about Asset U—than when the third asset is similar to Asset F. The results were stark: despite the choice and learning environment being held constant, people went from selecting the unlikely-outperforming Asset U less than 20% of the time when the third asset was similar to Asset F ("Interference (Asset F)") to selecting Asset U more than 60% of the time when the third asset was similar to Asset U ("Interference (Asset U)").[4] This more-than-40 percentage points shift in choices in the sequential information environment led to an even higher (revealed) preference for the unlikely-outperforming asset than the simultaneous information environment. People's beliefs matched the choice patterns: they were more optimistic about the performance of Asset U (F) when the third asset was more similar to Asset U (F).

Figure 2 summarizes results across conditions. Despite facing the same decision and having the same information, the interaction between the learning environment, attention, and memory shifted people's choices by nearly 50 percentage points. This dramatic difference in revealed preferences highlights the importance of understanding the role of cognition in driving choice behavior. The analogous data on beliefs shows that the standard assumption that people's perception of risk matches the objective distributions conveyed to them may not hold. Exploring the factors which affect mental representations is critical for understanding risky choice.

Our findings have implications for the observed "decision-experience" gap in risky choice (see Hertwig et al. (2004), and de Palma et al. (2014) for review). In prior work documenting this

[4]The shift in choices resembles a "decoy effect" Tversky and Simonson (1993) but is conceptually different because the manipulated attributes were payoff-irrelevant.

Figure 2: Learning about Uncertainty. Our Baseline Experiment documents a highly significant choice gap as a function of the learning environment: Most choose the unlikely-outperforming Asset U when learning from simultaneous information but switch to choosing the frequently-outperforming Asset F when learning from sequential information, despite both assets having the same marginal distributions. This gap can be eliminated and reversed by manipulating memory and attention. Having participants keep a record of the sample ("No Recall Bias"), reducing the length of the sample to minimize memory requirements ("Minimal Requirement"), and interfering with memory formation by introducing an attractive third asset that is similar to Asset U ("Interference (Asset U)") reduces the fraction of choices for Asset F after sequential information from 84% down to 32%. Meanwhile, minimizing the salience of the infrequent, extreme state ("No Salience Bias") increases the fraction of choices for Asset F after simultaneous information from 42% to 66%.



gap, people select between a safe option and a risky gamble that has a large but unlikely payout. In the description environment, information about the outcomes and their respective likelihoods is presented simultaneously. In the experience environment, people learn this information by sequentially sampling from the distribution of the gamble, with replacement; they can stop sampling whenever they want and then proceed to make a choice. People were more likely to select the risky asset in the former environment than the latter. Several explanations have been proposed for this shift in choices. In the case of learning from description, the higher preference for the risky asset has been attributed to probability weighting, which leads to the overweighting of unlikely events (Kahneman and Tversky, 1979). In the model, objective probabilities are transformed

through a probability-weighting function—typically assumed to be a primitive—which leads people to overweight low probability outcomes and underweight high probability outcomes (Tversky and Kahneman, 1992). Some have argued that the apparent underweighting of unlikely events when learning from experience is due to differences in information. For example, Fox and Hadar (2006) propose that people are more likely to go for the safe option because they do not even know about the unlikely payout of the gamble—they simply may not sample enough to have the same information about the distribution as those in the experience treatment (see Cubitt et al. (2022) for a similar conjecture). Others argue that the gap is driven by reinforcement learning (Erev and Barron, 2005) or simple tallying heuristics (Hills and Hertwig, 2010). Notably, Erev et al. (2010) do propose limited memory as one of many potential explanations, but do not test this mechanism directly.

We derive the "description-experience" gap from first principles—namely, bounded attention and memory—and present direct evidence for these mechanisms within the same paradigm. Importantly, we predict and show that the gap in choices is driven by a difference in perceptions of and beliefs about uncertainty. We do so in a setting that is explicitly designed to rule out the explanations that have been previously proposed in the literature. The experiment 1) involves a choice between assets with the same outcome distribution (which rules out primitives such as probability weighting), 2) ensures that people have the same information about the distribution across learning environments (which rules out sampling error), 3) abstracts from differences in hedonic feedback before the choice is made (which rules out reinforcement learning), and 4) allows for the explicit manipulation of factors affecting attention and memory (which allows for identification of the mechanism). The fact that we can eliminate and even reverse the "description-experience" gap by manipulating the relevant cognitive constraints suggests that the gap is not driven by something inherent to learning from description or experience per se.

Our results contribute to the literature on experience effects in consumer choice and financial decision making. Studies have documented long-lasting effects of personal experiences on economic decision-making; for example, Malmendier and Nagel (2011) find that negative stock-market experiences during the Great Depression lead to more conservative investing decades later. Malmendier and Wachter (2021) review this work and link it directly to memory. Other studies show that generating *artificial* experience through sequential sampling can be used to affect beliefs about returns

and investment choices (e.g., Kaufmann et al., 2013; Laudenbach et al., 2022). We contribute to this research on experience effects by highlighting the role of bounded memory as the essential link between beliefs and choices in sequential learning environments.

We also contribute to the more conventional portfolio choice and asset pricing literature by varying dependence while carefully holding marginal distributions constant. Dependence is intrinsic to models of optimal portfolio choice and capital market equilibrium going back to Markowitz (1952) and Sharpe (1964). In these models, dependence between asset returns is important because it results in non-diversifiable risk (systematic risk or 'beta'), which makes portfolios more risky for individuals and depresses prices in the aggregate. Going beyond the effect of dependence on portfolio risk, we provide the underpinnings of recent studies focusing on effects of dependence on beliefs, portfolio choice, and asset prices when investors learn from realistic samples of returns and prices (Ungeheuer and Weber, 2021, 2023). Incorporating dependence is unusual in behavioral finance, where typical approaches assume narrow framing when investors evaluate individual assets (e.g. Barberis et al., 2016, where investors evaluate stocks in isolation using the Cumulative Prospect Theory value function and a sample of 60 monthly returns). The few exceptions that explicitly consider dependence being Cosemans and Frehen (2021), who correlate historical stock returns with return predictors based on Salience Theory, and Barberis and Jin (2023) who consider the implications of reinforcement learning for investor behavior. We contribute to this literature by showing how dependence can affect capital markets through a fundamentally new interaction between learning environments, memory and attention mechanisms (not related to portfolio risk), and belief formation.

Finally, our findings inform the recent work exploring how cognitive constraints reconcile and rationalize seemingly anomalous economic phenomena in a variety of domains. Models of bottom-up attention and salience have been used to explain the Allais paradox and risk-seeking behavior (Bordalo et al., 2012), insensitivity to taxes (Chetty et al., 2009; Taubinsky and Rees-Jones, 2018), and behavior in strategic interactions (Li and Camerer, 2022); see Bordalo et al. (2022b) for a review. Models of bounded memory have been used to explain projection bias, context-dependent consumer choice, availability and representativeness heuristics (Bordalo et al., 2022a, 2020), as well as overreaction to information (Enke et al., 2020; Hartzmark et al., 2021), and belief-decay (Graeber et al., 2022). Recent studies build on evidence from neuroscience to model how the

difficulty of perceiving and adapting to outliers can explain the neglect of outlier risk (d'Acremont and Bossaerts, 2016; Payzan-LeNestour and Woodford, 2022). Finally, models of cognitive noise and complexity have been used to explain small-stakes risk aversion (Khaw et al., 2021), insensitivity to decision-relevant parameters (Enke and Graeber, 2023), the four-fold pattern of risk (Oprea, 2022), and non-constant time discounting (Enke et al., 2023). We contribute to this literature by identifying an important interaction between cognitive constraints and the learning environment, and documenting its effect on perceptions of uncertainty and choice behavior.

The rest of the paper proceeds as follows. In Section 2, we outline the framework and predictions on uncertainty perception and risk-taking as a function of the learning environment. We introduce our experimental design in Section 3. Section 4 presents the results of our experiments for both choices and beliefs, and discusses alternative explanations. We conclude in Section 5 by discussing potential implications of our findings.

2 Cognitive Constraints and the Learning Environment

We now discuss a framework of bounded memory and attention and present predictions on how these mechanisms interact with the learning environment to impact risky choice. Appendix A derives the predictions formally. This framework is tightly linked to our experimental design presented in Section 3, which was developed to allow for theory-driven manipulations of the relevant cognitive constraints.

In the framework, the decision-maker (DM) first recalls information about an asset's distribution of outcomes and then evaluates this information when making a choice. With perfect memory and no attentional distortions, the learning environment has no impact on choice—recalled information at the time of decision is the same regardless of whether outcomes are sampled sequentially or the distribution is provided simultaneously—and the DM makes a decision in line with Expected Utility Theory. However, allowing for imperfect recall can lead to behavior that seems to underweight unlikely outcomes, while attentional distortions will lead to the overweighting of outcomes in unlikely but salient states, conditional on being recalled.

Let each asset be denoted as A^i and each possible state as $s \in S$, which occurs with probability π_s. Let x_s^i correspond to the outcome of A^i in state s. The DM's subjective value from a given

outcome is denoted $v(x_s^i)$. We first consider the case of two assets, A^U and A^F. The DM with perfect memory and no attentional distortions will compute the following value $V^R(A^i)$ for each asset:

$$V^R(A^i) = \sum_{s \in S} \pi_s * v(x_s^i). \tag{1}$$

where the superscript R corresponds to the rational benchmark. A DM in our setting will depart from this baseline due to a) imperfect recall of outcomes because of memory constraints and b) over/underweighting likelihoods due to bottom-up attention. To preview the results, recall will be adversely affected by the number of states as well as the similarity of outcomes across states and assets; over/underweighting will be affected by the salience of states, which is a function of both the payout distribution and visual factors.

2.1 Gathering information

To make a choice, the DM assembles information about the distribution of risk either from being presented with it directly or by learning through sequential sampling. In the case of the latter, at the time of choice she must accrue this information through memory recall. The process of memory recall in our framework builds on the models of Bordalo et al. (2020) and Anderson et al. (1998). Let $S(u, v) : E \times E \to [0, 1]$ measure the similarity between two outcomes u and v in a DM's set of sampled outcomes E, where $S(u, v) = 1 \iff u = v$. Let χ_s^i be the remembered outcome in state s for asset A^i. The DM correctly recalls a given outcome (i.e., $\chi_s^i = x_s^i$) with probability $r(x_s^i)$, which is a function of similarity both within asset and across states, and within a state across assets. With probability $1 - r(x_s^i)$, the DM instead recalls the mean of the outcomes most similar to the recall target in that state. That is, the DM experiences interference in her ability to recall outcomes from two sources: the similarity of outcomes across assets within state and the similarity of outcomes across states within asset. Recall is perfect when outcomes are entirely dissimilar across both states and assets, and recall deteriorates as either type of similarity increases.

2.2 Forming mental representation

Once the DM has access to the distribution, either through the recall process described above or by being presented with it directly, she forms a mental representation of the outcomes and respective

likelihoods. This mental representation is subject to attentional distortions as a function of state-specific salience. Namely, states that are more (less) salient will attract more (less) attention and be over(under)weighted in the decision-making process.

A large literature in economics and psychology has shown the importance of attention in risky choice. Limited attention is channeled towards some states over others as a function of their salience (Brandstätter and Körner, 2014; Spitmaan et al., 2019; Bordalo et al., 2022b). As a result of channeled attention, states that are more salient are overweighted in the decision-making process compared to states that are less salient (Frydman and Mormann, 2016). In the context of risky choice, a state's salience is determined by both visual stimuli (Frydman and Mormann, 2016; Bose et al., 2022) and payoffs (Bordalo et al., 2013).[5] Notably, as shown in Bose et al. (2022), Li and Camerer (2022), and Bruce and Tsotsos (2009), amongst others, the salience of an object can be predicted using algorithms such as the Graph-Based Visual Saliency model (Harel et al., 2006). We use this technique to validate our assumptions on the salience ranking of states in our experimental studies.

When forming a mental representation, the weight assigned to a state is a function of its salience. Specifically, in the framework, states are first are ranked by their relative salience. Then, objective state probabilities are replaced with distorted probabilities, where more salient states are over-weighted and less salient states are underweighted. Let $\sigma(\chi_s^i, \chi_s^j)$ be a function mapping recalled outcomes to state salience and k_s be the relative ranking of state s by salience.[6] The weight then given to state s is:

$$\omega_s = \frac{\delta^{k_s}}{\sum_{r \in S} \delta^{k_r} * \pi_r}, \tag{2}$$

where $\delta \in (0, 1]$. Finally, the DM evaluates each asset by computing:

$$V(A^i) = \sum_{s \in S} \hat{\pi}_s^i * v(\chi_s^i), \tag{3}$$

where $\hat{\pi}_s^i = \pi_s^i * \omega_s$.

[5] For empirical evidence on the importance of state-based salience in the domain of belief-updating, see Ba et al. (2022).

[6] See Appendix A for the specific properties of σ.

2.3 Choice behavior

Consider a decision environment with a choice between two assets A^U and A^F and a finite number of equally-likely states $|S|$, where $x_s^F > x_s^U$ for all but one state s^*, where $x_s^F < x_s^U$. Finally, as in our experimental setting, for each outcome in A^U, there exists an equivalent outcome in some state for A^F and vice-versa. A key property of this environment is that $V^R(A^U) = V^R(A^F)$. Our first prediction is on the choice between assets under perfect recall.

Prediction 1. *With no memory constraints, a DM with salience-distorted attention will prefer A^U to A^F.*

This prediction is relevant to the "simultaneous" learning environment where the DM is presented with information about the potential outcomes simultaneously. There, imperfect recall does not play a role since all of the information necessary to form an accurate mental representation is present at the time of choice. Our second prediction introduces the possibility of imperfect recall of the relevant outcomes.

Prediction 2. *Consider a setting with imperfect recall. Suppose the similarity between outcomes is nonzero across both states and assets. Then there exists a number of states $|S^*|$, such that for $|S| > |S^*|$, the DM will prefer A^F to A^U.*

The intuition for the second prediction is straightforward. If the DM misremembers the state in which $x_s^U > x_s^F$ while remembering any number of states where $x_s^U < x_s^F$, then she will prefer A^F to A^U. Consider the structure of A^U and A^F, where the former outperforms the latter in one state while the latter outperforms the former in each of the others. This implies that as the number of states grows, the chance that the DM will misremember the state in which $x_s^U > x_s^F$ while remembering a state where $x_s^U < x_s^F$ will increase as well. Since imperfect recall plays a bigger role in the "sequential" learning environment than in the "simultaneous" one, the first two predictions imply a preference reversal: A^F will be selected over A^U in the former and A^F will be selected over A^U in the latter.

Our third prediction is on how similarity impacts choice behavior through the recall process.

Prediction 3. *Consider a setting with imperfect recall and a total number of states $|S|$. The likelihood that A^F is preferred to A^U is increasing in the similarity of outcomes both across assets*

within a state and across states within an asset.

The intuition for the third prediction follows from the second. Similarity between outcomes increases the chance that they will be misremembered. The relative impact of similarity is higher for the outcomes of infrequent states simply because there are less of them. Thus, conditional on a number of states $|S|$, the DM is (weakly) more likely to prefer A^F to A^U as the similarity between outcomes increases.

We build on this intuition for our last prediction, which directly tests the interference aspect of imperfect recall. Consider a third asset A^D that is better than both A^U and A^F, which the DM learns about but cannot choose. Interference predicts that one could reverse the preference between assets in the "sequential" environment when the third asset is more similar to asset A^U than asset A^F.

Prediction 4. *Suppose a third asset A^D is included, which dominates both A^U and A^F in every state. The likelihood that A^U is preferred to A^F is increasing in the similarity between the outcomes of A^D and A^U.*

The logic for this prediction comes from the property of interference that underlies most models of memory. When an outcome from a state is misremembered, the outcome that replaces it is more likely to be associated with a similar source. If the "replacement" outcomes from the similar source are superior to those of both assets in the choice set, then misremembering an outcome from A^U will increase its attractiveness. Including the similar asset also increases the chances that an outcome from A^U is misremembered in the first place. Both forces lead to an increased preference for A^U in the sequential learning environment.

Notably, as we show in Appendix A, including a third asset A^D that is more similar to A^F maintains the preference for A^F over A^U in the sequential learning environment.

3 Experimental Design

In this section, we first describe the essential features of our Baseline paradigm that is used to study the impact of the learning environment on risky choice.[7] Participants ($N = 785$) were

[7]The experiment was pre-registered on AsPredicted.org, see this link. Complete instructions are reported in a supplementary document provided under this link.

recruited using the Prolific crowdsourcing platform (Gupta et al., 2021) and restricted to investors from the US and UK. Summary statistics for participants are reported in Table IA1 in the Internet Appendix.

The experiment's overall structure is illustrated in Figure IA1. Participants first received introductory instructions. These included descriptions of the upcoming tasks, the information they would have access to, as well as an explanation of how they would be paid. Participants then went through two rounds of investment decisions. In each, they saw information about two risky assets. As outlined below, participants learned about each asset's distribution by either being presented with the outcomes simultaneously or by sequentially sampling them. After being given this information, participants chose between the two assets and stated their beliefs about the assets' outcomes. Beliefs about risk were elicited in the form of probabilities for an outcome less than \$1, equal to \$1, or greater than \$1. Both assets' actual distributions' were discrete uniform between \$0 and \$2, so that \$1 is always the mean and median.[8] For the choice between assets, we randomize the order of presentation both within-participants across rounds and across participants. To account for potential spill-over effects, half of participants were first asked for their choice and incentivized accordingly ("choice group"), and the other half was first asked for their beliefs and incentivized for accuracy ("beliefs group"). The experiment concluded after the second round with some final questions (e.g., self-assessed risk aversion) and participants were informed about their bonus payment, which depended on their incentivized task.

One of the two rounds was randomly selected for the bonus payment at the end of the experiment. The "choice group" was later paid with a bonus drawn from the selected asset's distribution (outcome between \$0 and \$2, on average \$1). The "beliefs group" was paid for a randomly selected belief (e.g., the average of Asset F's outcomes) with a bonus of \$2 if their stated belief is within $+/-0.03$ of the statistically correct value. Including a base fee, the average hourly salary for our Baseline Experiment was more than \$13.

[8] Danz et al. (2022) show that the binarized scoring rule results in bias in elicited beliefs and higher error rates. They argue that incentivizing truthful reporting using belief quantiles, which we use here, results in more truthful reporting.

3.1 Varying the Learning Environment

The joint distribution of the two risky assets is described in Figure 1. Both assets exhibited precisely the same marginal outcomes, uniformly distributed between \$0 and \$2.[9] However, the structure of dependence is asymmetric. Asset F moderately outperformed Asset U most of the time, while underperforming by a large margin once.[10] In the Baseline study, participant made choices between assets with a "short" distribution consisting of 11 states (or "situations") in one round, as shown in Figure 1, as well as a "long" distribution with 51 states in another round. The order of the two rounds was counterbalanced, i.e., some participants first saw the short and then the long distribution, or vice versa.

Participants learned about the joint distribution of the assets in two different environments. These environments varied between-subject and represent our main treatments. Half the participants were given information that described the outcomes of both assets in all states (or "situations"), along with the probability of each state. This constituted our Simultaneous Information treatment. The other half learned about the two risky assets by sequentially sampling each potential state and seeing the outcomes from the joint distribution. Each was told that 1) they would see the outcomes of each potential state and 2) that each state was equally likely to be realized. This constituted our Sequential Information treatment. It is important to note that presenting participants with every state in the Sequential Information treatment ensured that they were exposed to the same information about the relevant distributions as those in the Simultaneous Information treatment at the time of choice.

Figure 1 illustrates the Simultaneous Information learning environment. The graph shows the outcome of both assets in all states. In the instructions before the first round, participants also received an explanation of the graph and the information that all states are equally likely. They needed to pass a comprehension test to continue to the first round. Underneath the graph, a verbal statement also describes the distribution, as is common in traditional decision-making experiments. The statement mentions all joint outcomes and their probabilities, thus fully describing the joint outcome distribution.[11] In the Simultaneous Information treatment, participants could also view

[9] This design was inspired by the paradigm used in Paterson and Diekmann (1988) to study violations of Expected Utility axioms.

[10] The participants' instructions referred to Asset U and F as Asset A and B, respectively.

[11] See Figure IA2 for the full Simultaneous Information environment including the verbal statement. The graph

the graph when making the choice between assets and reporting beliefs, removing the need to memorize the distribution.

Figure 3 illustrates the Sequential Information learning environment, which uses both graphical (bar chart) and numerical information (realizations in $) to clearly present all outcome pairs.[12] Participants drew 11 (51) outcome pairs in the short (long) round of the experiment. Each sample is fully representative—presenting participants with the outcomes in every state—so that there is no sampling error. Additionally, each sample is randomly ordered, so that particular outcome orders cannot explain systematic treatment effects. Random ordering also allows us to analyze the importance of recency and primacy effects, which we explore in Section D.1.4. Participants chose how long to view the outcomes of each state and when to move on by clicking "Next." There was no option to go back. We collected viewing times for each state, which we use to analyze attention effects in Section 4.3. In the Sequential Information treatment, participants could not view a "simultaneous" summary of the relevant distributions while making their asset choices and reporting their beliefs.

Figure 3: Baseline Experiment – Sequential Information Treatment.



was designed to highlight the differences between assets using both visual presentation and text. Follow up treatments which remove both the text and box from the graph, while keeping the state ordering constant, yielded the same results as this Baseline Simultaneous Information treatment.

[12]This sampling design was inspired by the experiment in Ungeheuer and Weber (2023) who study the impact of outperformance on investment choice.

3.2 Discussion of Design

There are several noteworthy aspects of the Baseline experimental design. First, showing a a fully representative sample in the Sequential Information learning environment ensures that participants have been exposed to the same information at the time of choice, which accounts for the oft-cited sampling error explanation for the description-experience choice gap (Fox and Hadar, 2006; Cubitt et al., 2022). Second, presenting the state-by-state outcomes in the Simultaneous Information learning environment accounts for the "mere-presentation effect" as a potential explanation (Erev et al., 2008).[13] The same design choice also accounts for the "unpacking-and-repacking" explanation (de Palma et al., 2014) which also relies on a difference in emphasis on unlikely outcomes between the Simultaneous and Sequential Information learning environments. Third, we used assets with the same underlying distributions to rule out explanations based on primitives such as probability weighting, since models like original or cumulative prospect theory (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992) would not predict a preference for either asset in the Simultaneous Information environment. Finding a strict preference for the unlikely-outperforming Asset U in the Simultaneous Information environment would thus provide evidence for bottom-up attention as a driver for the overweighting of unlikely events. Finally, employing a sequential learning environment where outcomes were not payoff relevant accounts for reinforcement learning as an explanation for any observed differences in choices (Erev and Barron, 2005).

Our design differs from prior work on risky choice by including a graphical depiction of state-by-state distributions of outcomes. We did this for several reasons. First, the graphical depiction ensures that the presentation of events is proportional to their probability, as is the case in the Sequential Information learning environment. Second, as outlined further below, the visual representation allows us to manipulate multiple aspects of salience exogenously and to then validate those manipulations using existing tools from cognitive psychology.

[13] Erev et al. (2008) argue that the description-experience choice gap is due to the typical presentation of lotteries in the description learning environment, e.g., $((\$x, p; \$0, 1-p) = (\$8, 0.1; \$0, 0.9)$, affording propositional representation to the unlikely events. The same lottery's presentation through sampling is afforded analogical representation, where the unlikely outcome is also unlikely to be presented. Our design accounts for this by presenting the outcome of each state individually both in the simultaneous and Sequential Information learning environments.

3.3 Testing the Mechanism

We designed five additional treatments to directly test the proposed framework by exogenously manipulating memory and attention. All variations were theory-driven, motivated by the model outlined in Section 2 and Appendix A. The framework predicts that manipulating salience can drive choices away from the unlikely-outperforming Asset U towards the frequently-outperforming Asset F in the Simultaneous Information environment. It also predicts that manipulating memory constraints can decrease the preference for the frequently-outperforming asset in the Sequential Information environment, and reverse the preference altogether.

3.3.1 Manipulating Salience

We sought to manipulate salience without changing the state-by-state distribution of the assets. To do so, we leveraged the Attention based on Information Maximization (AIM) model of Bruce and Tsotsos (2009) to change the presentation of the states from one where the salient state clearly deviated from the local environment—and thus attracted attention—to a depiction where this deviation was less unexpected.[14] We did this by changing the order of the states from the predictably-increasing pattern in the Baseline treatment—where differences in outcomes in adjacent states are relatively small except for where Asset U outperforms—to a randomized order where adjacent states could have large differences in the levels of outcomes—where a large outcome difference within a state would be less surprising. The information presented to participants ($N = 100$) in this No Salience Bias treatment is depicted in Figure 4 below; otherwise the design was identical to the Baseline Simultaneous Information treatment.[15] It is important to note that the underlying distributions and the numbers associated with each state were the same across the Baseline and No Salience Bias treatments.

To test our assumption that the No Salience Bias treatment decreased the salience of the state where the unlikely-outperforming asset was superior (state 11), we used the bottom-up Graph-

[14] In the AIM model, salience emerges from first principles of bounded rationality—efficient coding and information theory. The salience of some component (in our case, the state) is a function of the local information relative to the surrounding information; a state is more salient if the content is unexpected given the surrounding context. See Appendix A for further discussion.

[15] We still display the verbal description of all outcome pairs and probabilities, as in the Baseline Simultaneous Information treatment.

Based Visual Saliency model and the associated algorithm proposed in Harel et al. (2006).[16] The output of the model for both treatments is presented in the bottom row of Figure 4. As predicted, in the Baseline Simultaneous Information treatment, the outcomes associated with the state where Asset U outperforms are by far the most salient; the large contrast between payoffs relative to the local environment is predicted to channel more attention to that state than other states. However, changing the presentation such that the differences in payoffs within a state are less surprising substantially decreased the salience of state 11. No state is clearly dominant in terms of salience in the No Salience Bias treatment, implying that salience is less likely to bias choice towards a particular asset in this setting. Based on Prediction 1, people should choose Asset U less in the Random Order treatment than in the Baseline Simultaneous Information treatment.

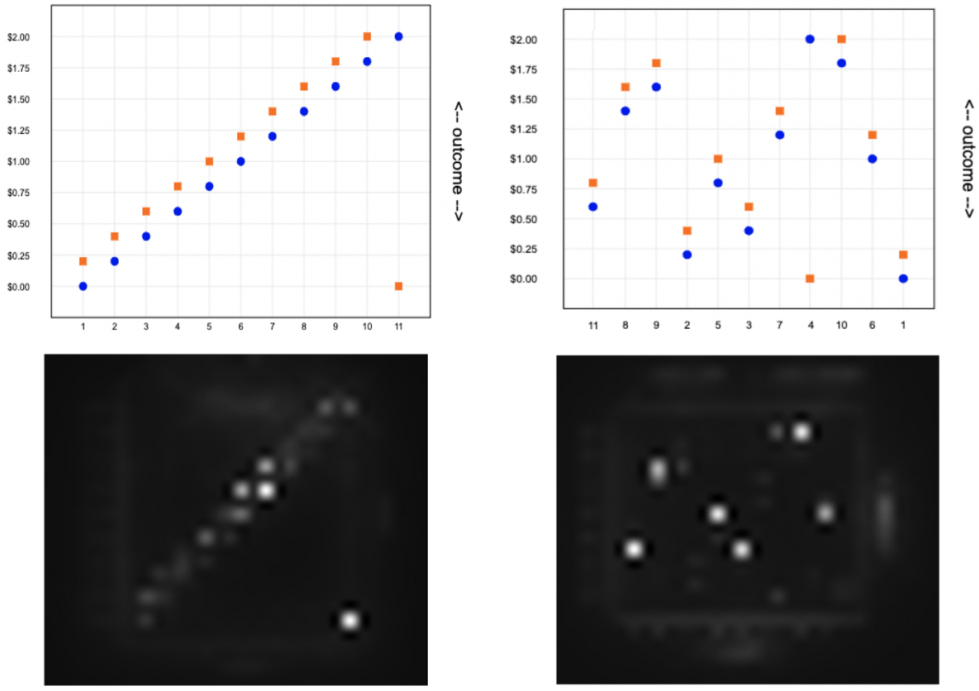
3.3.2 Manipulating Memory

Our first two memory manipulations sought to decrease the impact of imperfect recall on choice. The No Recall Bias treatment proceeded in the same way as the Baseline Sequential Information treatment, except that participants ($N = 100$) were asked to write down outcomes during the sequential sampling process. Specifically, they were asked to have a pen and paper (or other method of recording) ready and incentivized with a potential bonus payment for correctly responding to a quiz at the end of the round (see Figure IA3). By reducing the scope for imperfect recall, participants in the No Recall Bias treatment are predicted to display a smaller preference for Asset U than those in the Baseline Sequential Information treatment.

Based on the comparative static in Prediction 2, the Minimal Requirement treatment reduced the number of states to potentially remember from 11 to 3—a reduction below critical thresholds for memory capacity (e.g., Miller, 1956). This represents the minimal example where it is still possible to have "frequent outperformance" of Asset F (in 2 of 3 states) while maintaining the same marginal distributions for both assets ($0, $1, and $2 with equal likelihood). Reducing the number of states decreases the likelihood that 1) the state where Asset U outperforms is misremembered while 2) a state where Asset F outperforms is correctly remembered. In turn, participants ($N = 99$) in the Minimal Requirement condition are predicted to prefer Asset U to a smaller extent than those in

[16] In short, the model generates a saliency map of an image using an algorithm that was validated through measures of human visual fixation. The output of the saliency map is based on levels of activation in units of pixels—higher pixels/brighter areas are associated with higher levels of salience.

Figure 4: This figure depicts the relative salience of the unlikely-outperforming state of Asset U in the Baseline simultaneous information (left) versus No Salience Bias simultaneous information (right) treatments. The top two graphs correspond to the distributions shown to participants. The bottom two figures correspond to the salience maps of those distributions as measured by the Graph-Based Visual Saliency Method. In the Baseline simultaneous information treatment, the performance of Asset F in state 11 was by far the most salient area, with an average activation level of approximately 230 pixels (the next highest outcome had an activation level of less than 100). In contrast, in the No Salience Bias simultaneous information treatment, the outcomes associated with state 11 had an average activation level of approximately 30 pixels, while six other states had outcomes in the range of 100 to 150 pixels.



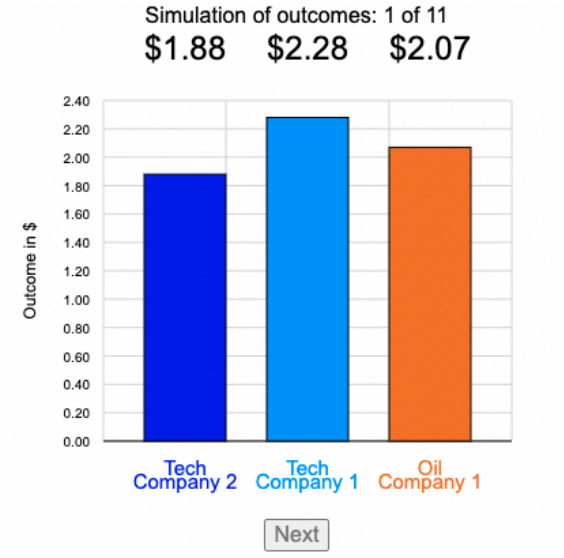
the Baseline Sequential Information treatment.

Our third manipulation sought to test Prediction 4 directly. Here, we exploited the role of interference in an attempt to manipulate memory formation in the Sequential Information learning environment. To that end, we introduced a decoy Asset D which could never be chosen; the choice remained the same—between Asset U and F. Participants sequentially sampled outcomes from Asset D along with the two other assets.[17] The decoy asset's outcomes dominate the outcomes of the two assets that can be chosen, so that it should be perceived as very attractive. Prediction 4 says that assets which are similar to the decoy should also be perceived as more attractive due to interference in memory formation, even if the features that create similarity are payoff-irrelevant.

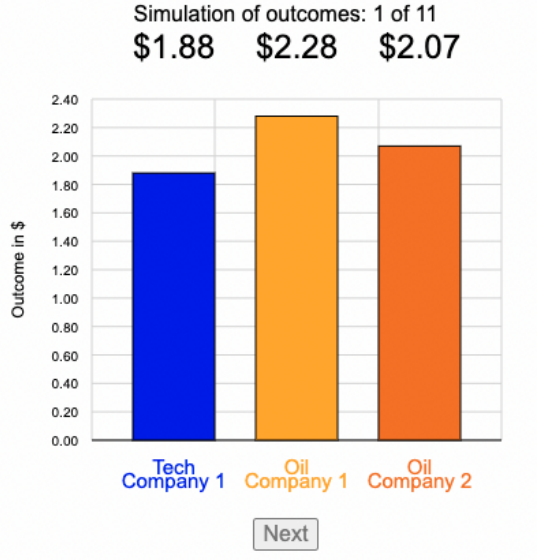
[17] See Table IA2 for the full joint distribution.

Using color (shades of orange vs. blue) and industry classification (oil vs. tech), we structured the decoy Asset D to either be similar to Asset U or F. Figure 5 shows the conditions where Asset D was similar to Asset U ("Interference (Unlikely-Outperformance)" condition) and where Asset D was similar to Asset F ("Interference (Frequent-Outperformance)" condition). Participants were randomized into the former ($N = 99$) or the latter condition ($N = 102$). While participants in the Interference (Frequent-Outperformance) condition are predicted to continue selecting Asset F over U, the framework predicts that not only will the gap decrease in the Interference (Unlikely-Outperformance) condition, but it may actually reverse.

Figure 5: This figure illustrates the two interference conditions. Participants sequentially sampled an additional Asset D but could only select between the other two assets, which exhibit the same marginal distributions as Assets U and F in the Baseline design. Asset D dominates the two assets that can be chosen. In the 'Unlikely-Outperformance' variation (a), the dominant Asset D is similar (in color and industry) to the unlikely-outperforming Asset U ('Tech Company 2'). In the 'Frequent-Outperformance' variation (b), Asset D is similar to the frequently-better Asset F instead ('Oil Company 2'). Note that the ordering of assets was randomized.



(a) Interference (Unlikely-Outperformance)



(b) Interference (Frequent-Outperformance)

4 Empirical Results

This section reports the results on choices (Section 4.1) and beliefs (Section 4.2). We explore additional mechanisms in Section 4.3. We start with our Baseline Experiment and then proceed

with the theory-testing conditions.[18]

4.1 Choices

Figure 2 presents results across the Baseline and theory-testing conditions. Table 1 shows our results on choices across the learning environments, with the Baseline conditions marked by E1 and D1.

Table 1: Main Results – Choices. In this table, we report average choices. We report the average fraction of participants selecting each of the alternatives. t-statistics are based on participant-clustered standard errors whenever both rounds are included. */**/*** indicates statistical significance at the 10%/5%/1% level.

Condition	Choice Asset U	Choice Asset F	Difference versus 50%	t-test	N
Simultaneous Information					
(D1) Baseline	58.01%	41.99%	-8.01%***	-3.16	381
(D2) No Salience Bias	34.00%	66.00%	16.00%***	3.36	100
Sequential Information					
(E1) Baseline	15.10%	84.90%	34.90%***	19.57	404
(E2) No Recall Bias	37.00%	63.00%	13.00%***	2.68	100
(E3) Minimal Requirement	37.37%	62.63%	12.63%***	2.58	99
(E4.1) Interference (Frequently-Outperforming)	19.61%	80.39%	30.39%***	7.69	102
(E4.2) Interference (Unlikely-Outperforming)	67.68%	32.32%	-17.68%***	-3.74	99
Difference Across Learning Environments					
(E1-D1) Baseline			42.91%***	13.98	785
(E2-D2) Eliminating the Gap			-3.00%	-0.44	200

Consistent with our first Prediction 1, people learning about the distribution in the Simultaneous Information condition (D1) had a significant preference for the *unlikely-outperforming* Asset U, choosing it 58.01% of the time ($p < .01$).[19] On the other hand, consistent with Prediction 2, learning the same information through sequential sampling (E1) shifted choices towards the *frequently-outperforming* Asset F, which was selected 84.90% of the time ($p < .01$).

The gap in choices between the Simultaneous and Sequential Information environments (E1-D1) is substantial, at 42.91 percentage points ($p < .01$). The existence of the gap is surprising given that the underlying distributions and information about the assets were kept constant, which rules out many potential mechanisms such as prospect theory or sampling error.

[18] We present the 11-state version of the experiment here. As there were no significant differences in the order in which beliefs and choices were elicited, we pool across these variations when presenting our main results. The Internet Appendix contains additional results and sample splits.

[19] Unless otherwise noted, p-values correspond to the difference from 50%.

We next vary the Baseline Experiment's design to exogenously manipulate memory and attention. Decreasing the salience of the state where Asset U outperforms (condition D2) flips the preference for the unlikely-outperforming asset towards the frequently-outperforming: the fraction of participants selecting Asset F increases from 41.99% to 63.00% ($p < .01$). The elimination—and even a reversal—of the preference for the asset with the unlikely upside suggests that attention and salience play an integral role in driving decisions from Simultaneous Information.

Similar to decreasing salience bias with treatment D2, we sought to reduce the bias induced by imperfect recall in the sequential learning environment with the No Recall Bias treatment. Row (E2) of Table 1 shows that decreasing the scope for imperfect recall reduced the fraction of choices for the frequently-outperforming Asset F from 84.90% to 63.00% ($p < .01$ both between treatments and relative to 50%). Notably, comparing the No Salience Bias and No Recall Bias treatments (E2-D2) shows the complete elimination of the choice gap between learning environments: people's choices were essentially the same regardless of whether they learned from simultaneous (D2) or sequential information (E2). This highlights the role of attention and memory constraints as the drivers of the gap in the Baseline conditions, rather than something inherent to the learning environment per se.

Next, as a separate test of the imperfect recall channel, we reduced the number of states from 11 to 3 in the Minimal Requirement treatment. According to Prediction 2, the lower number of states that need to be recalled should reduce the memory bias, and in turn, the preference for the frequently outperforming Asset F. Row (S3) of Table 1 shows that this is indeed the case: Asset F was chosen more than 20% less in the Minimal Requirement treatment than in the Baseline treatment (62.63%; $p < .01$ both between treatments and relative to 50%).

Finally, we move to testing the memory mechanism by manipulating interference in the sequential learning environment (Prediction 4). When the decoy asset is more similar to the frequently-outperforming asset in the Interference (Frequent-Outperformance) treatment, this merely reinforces the preference for Asset F. Indeed, as shown in row (E4.1) of Table 1, people select this asset at around the same rate as the the Baseline Sequential Information condition (E1). However, this preference changes dramatically when the decoy asset is more similar to the unlikely-outperforming asset. As shown in row (E4.2), the preference for Asset U goes from 15.10% in the Baseline Sequen-

tial Information condition to 62.89% in the Interference (Unlikely-Outperformance) condition.[20]

Adding a decoy asset that differed on non-payoff relevant attributes shifted choices from the vast majority of people preferring the frequently-outperforming Asset F to most people selecting the unlikely-outpeforming Asset U—a preference reversal of more than 40 percentage points! Manipulating similarity and interference led people to prefer the asset with *unlikely* outperformance in the sequential learning environment.

4.2 Beliefs

We now turn to examining beliefs. In line with the conjecture that the learning environment affects choices by changing people's perceptions of uncertainty, we find that beliefs largely follow the patterns in Section 4.1. Table 2 shows our main results on beliefs after learning in the simultaneous versus sequential learning environments. Beliefs about average outcomes are in Panel A and beliefs about the probability of outcomes above versus below $1 are in Panel B. Overall, we find that participants' beliefs are consistent with their choices. Note that because the order of the belief and choice questions were randomized, these results are not driven by spillover effects, e.g., a desire to report beliefs that are consistent with one's choices.

In our Baseline condition, participants become highly over-optimistic about the performance of the frequently-outperforming asset after learning through sequential information (E1). The distribution they sampled was the same for both assets, with an average (and median) of exactly $1. However, participants believed the average outcome is around $1.17 for the frequently-outperforming Asset F, which is more than $0.16 higher than their stated belief for the unlikely-outperforming Asset U ($p < .01$).[21] Consistently, they believe outcomes below $1 are less likely (likelier) than outcomes above $1 for the frequently-outperforming (unlikely-outperforming) asset, although all of these are the same (5 out of 11 situations below, 5 above $1). The difference-in-differences measuring the overoptimism for the frequently-outperforming Asset F versus the unlikely-outperforming Asset U is 25.87 percentage points ($p < .01$). On the other hand, participants were more optimistic about the performance of the unlikely-outperforming Asset U in the Simultaneous Information treatment (D1). They expected the Asset U to return an average of $0.07 more than Asset F

[20] A simple dislike for a specific color or industry cannot explain this result, as it would suggest an equal reduction of the fraction of choices for Asset F in both (E4.1) and (E4.2).

[21] Unless otherwise noted, p-values correspond to the difference between assets.

Table 2: Main Results – Beliefs. In this table, we report average beliefs. Panel A reports the average outcome expected for each of the alternatives (both marginal distributions have an expected value of \$1.00). Panel B reports the average probability expected for below/above-median outcomes for each of the alternatives (both marginal distributions have a 5-in-11 chance for these outcomes). The t-statistics in Panel B are from a difference-in-difference estimate between the below- and above-median probability assessments. t-statistics are based on participant-clustered standard errors whenever both rounds are included. */**/*** indicates statistical significance at the 10%/5%/1% level.

Condition	Beliefs Asset U	Beliefs Asset F	Difference (-in-Diff.)	t-test	N
Panel A: Expected outcome (0-2)					
Simultaneous Information					
(D1) Baseline	\$1.07	\$1.00	\$-0.07***	-2.83	381
(D2) No Salience Bias	\$1.06	\$1.09	\$0.03	0.81	100
Sequential Information					
(E1) Baseline	\$1.00	\$1.17	\$0.16***	7.34	404
(E2) No Recall Bias	\$1.08	\$1.08	\$0.00	0.06	100
(E3) Minimal Requirement	\$1.09	\$1.25	\$0.16***	3.20	99
(E4.1) Interference (Frequently-Outperforming)	\$0.99	\$1.38	\$0.38***	7.77	102
(E4.2) Interference (Unlikely-Outperforming)	\$1.20	\$1.08	\$-0.12**	-1.97	99
Difference Across Learning Environments					
(E1-D1) Baseline			\$0.23***	7.11	785
(E2-D2) Eliminating the Gap			\$-0.03	-0.66	200
Panel B: Risk assessment (below/above-median probabilities)					
Simultaneous Information					
(D1) Baseline – P($x < 1$)	39.09%	44.27%	6.46%***	3.46	381
– P($x > 1$)	42.63%	41.35%			
(D2) No Salience Bias – P($x < 1$)	41.71%	42.39%	3.26%	1.14	100
– P($x > 1$)	44.16%	41.58%			
Sequential Information					
(E1) Baseline – P($x < 1$)	48.40%	38.24%	-25.87%***	-12.58	404
– P($x > 1$)	35.04%	50.76%			
(E2) No Recall Bias– P($x < 1$)	38.08%	40.55%	-0.68%	-0.20	100
– P($x > 1$)	37.87%	41.02%			
(E3) Minimal Requirement – P($x < 1$)	35.87%	31.82%	-9.39%**	-2.32	99
– P($x > 1$)	32.08%	37.42%			
(E4.1) Interference (Frequently-Outperforming) – P($x < 1$)	47.98%	37.38%	-39.47%***	-8.05	102
– P($x > 1$)	28.83%	57.70%			
(E4.2) Interference (Unlikely-Outperforming) – P($x < 1$)	34.88%	49.67%	13.66%**	2.54	99
– P($x > 1$)	42.35%	43.48%			
Difference Across Learning Environments					
(E1-D1) Baseline			-32.34%***	-11.60	785
(E2-D2) Eliminating the Gap			-3.94%	-0.87	200

($p < .01$) and exhibit surplus probability mass of 6.46 percentage points for better outcomes of the unlikely-outperforming asset ($p < .01$).

There is a highly significant beliefs gap across the learning environments in our Baseline Experiment (E1-D1). Table 2's Panel A shows that this gap is \$0.23 for estimates of average outcomes ($p < .01$), i.e., participants expected an average outperformance of the frequently-outperforming asset over the unlikely-outperforming asset that is \$0.23 higher after learning from sequential information than from simultaneous information. Similarly, Panel B documents an excess probability mass of 32.34 percentage points ($p < .01$) for good versus bad outcomes of the frequently-outperforming asset relative to the unlikely-outperforming asset. To sum, the results on beliefs suggest that the systematic preference reversal in choices from Table 1 appears to be driven by a systematic change in perceptions—or mental representations—of uncertainty.

When we manipulate memory and attention, the gap in beliefs closes and then flips. We illustrate this for beliefs about average outcomes in (D2), (E2), and (E4.2).

First, decreasing salience bias in the Simultaneous Information treatment using the No Salience Bias condition led to more optimistic beliefs about Asset F by \$0.03, though the difference was not significant. Similarly, decreasing recall bias in the Sequential Information treatment with the No Recall Bias condition led to no significant differences in beliefs across the two assets. Attenuating biases in attention and memory closed the gap in beliefs as well as choices.

In the Interference (Unlikely-Outperformance) treatment, the decoy is predicted to interfere with recall about the unlikely-outperforming Asset U. In turn, participants were overoptimistic about Asset U (\$1.21) relative to Asset F (\$1.06) by \$0.14 ($p < .01$). This is a substantial reversal compared to the Baseline Sequential Information treatment where they were overoptimistic about Asset F compared to Asset U. The gap in beliefs flipped from \$0.23 ($p < .01$) in our Baseline conditions to \$-0.14 ($p < .05$) when we manipulate memory, even though the learning environment was kept constant.

4.3 Alternative Explanations

Despite presenting people with the same information in both the simultaneous and sequential learning environments, it may still be possible that people just did not attend to the state where Asset U outperformed in the latter. This would generate a gap in choices and beliefs through a

mechanism similar to *sampling error* (Fox and Hadar, 2006) or *selective sampling* (de Palma et al., 2014). We analyze viewing times for each outcome pair as a measure of attention to explore this possibility in our Baseline Experiment.

First, we find that participants do pay significantly more attention to extreme outcome pairs—consistent with salience-driven attention. The average number of seconds participants take to view a moderate outcome pair is 1.87, whereas states with extreme differences receive 3.65 seconds of viewing time, or 1.78 seconds more attention on average ($p < .01$; , see Table IA3). This near-doubling of attention for extreme outcome pairs confirms that the associated states are salient and not overlooked when people learn from sequential information.

Second, we test whether increased viewing times for states with extreme outcome differences are reflected in choices and beliefs. As shown in Figure IA4, participants who look longer at extreme outcome pairs compared to the average moderate outcome pair do not exhibit a weaker or stronger tendency to choose the frequently-outperforming Asset F compared to participants who do not pay more attention to these states. Table IA4 reports a regression showing that attention-weighted sampled outcome levels do not significantly predict choices. In summary, it seems that while salience does drive attention to states with extreme differences in the sequential learning environment, the bias generated by imperfect recall—as outlined in Prediction 2—plays the dominant role in affecting choices and beliefs.

5 Discussion

We find that the way that people learn about uncertainty changes their perception of it. People who learn from simultaneous information tend to overweight unlikely but salient outcomes while those who learn the same information sequentially tend to underweight such outcomes. These differences in mental representations are reflected in choices: learning about outcomes from "simultaneous" information leads to a preference for unlikely-outperforming assets, while learning the same information sequentially leads to a preference for frequently-outperforming assets. In line with our proposed framework, we demonstrate that the observed differences in beliefs and choices are driven by memory and attention constraints. This shows that the previously-documented "description-exprience" gap in risky choice can be derived from first principles of bounded ratio-

nality; manipulating factors related to cognitive constraints can close the gap completely, or even reverse it.

Our findings highlight how the same exact information can lead to vastly different perceptions of uncertainty depending on how it is learned. The kinds of unlikely but salient states that characterize the distributions in our experiment are not unusual, but rather common characteristics in a variety of settings. Fat tails and joint outliers are ubiquitous in a strongly connected society. The outsized influence of such outliers on individual and societal outcomes makes a better understanding of under- versus over-weighting of extreme states even more important.

Learning from sequential personal experiences has already been linked to choices of consumers and investors (Malmendier and Wachter, 2021). Our findings link this experience literature to a broader framework, which allows for a richer set of predictions on how behavior changes as a function of the learning environment and cognitive constraints. Integrating learning from sequential experiences into a common framework is also important because while the majority of studies have used "simultaneous" information about uncertainty—and developed models to fit the observed behavior (e.g., Prospect Theory)—learning from sequential information is an ecologically common way of learning about uncertainty.

Incorporating bounded attention and memory is also important because it helps to explain heterogeneity in risky choice (e.g., Diecidue et al., 2020). In a given learning environment, different subsets of participants may (i) correctly recall unlikely states and overweight them due to their salience, (ii) mis-remember unlikely states and neglect them (as if they "maximize the probability to be ahead"), (iii) or neither overweight nor underweight them (if the cognitive constraints are not binding).

Lastly, our results point to ways in which the learning environment can be tailored to improve the accuracy of mental representations of uncertainty. In our experiments, both the sequential and simultaneous learning environments led to biased beliefs. Exploring presentation formats that lead to more accurate beliefs is a promising avenue for research. Experimental research in the investment context already shows encouraging results. Kaufmann et al. (2013) document that a sampling procedure can be used to debias beliefs about returns of a single risky asset. In a context with multiple risky assets, Laudenbach et al. (2022) find that sequentially sampling return pairs instead of directly receiving simultaneous information about the joint return distribution improves beliefs

about dependence and makes portfolio choice more consistent with normative models (Markowitz, 1952). Our framework on how memory and attention affect choices could help structure the search for effective learning environments tailored to specific contexts (see the literature on nudging and boosting, e.g. Hertwig and Grüne-Yanoff, 2017).

References

Anderson, John R, Dan Bothell, Christian Lebiere, and Michael Matessa, 1998, An integrated theory of list memory, *Journal of Memory and Language* 38, 341–380.

Ba, Cuimin, J Aislinn Bohren, and Alex Imas, 2022, Over-and underreaction to information, *Available at SSRN* .

Barberis, Nicholas, Abhiroop Mukherjee, and Baolian Wang, 2016, Prospect theory and stock returns: An empirical test, in *Review of Financial Studies*, volume 29, 3068–3107 (Oxford University Press).

Barberis, Nicholas C, and Lawrence J Jin, 2023, Model-free and model-based learning as joint drivers of investor behavior, Technical report, National Bureau of Economic Research.

Bordalo, Pedro, John J Conlon, Nicola Gennaioli, Spencer Y Kwon, and Andrei Shleifer, 2022a, Memory and Probability, *The Quarterly Journal of Economics* .

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2012, Salience theory of choice under risk, *Quarterly Journal of Economics* 127, 1243–1285.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2013, Salience and Asset Prices, *American Economic Review* 103, 623–628.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2020, Memory, attention, and choice, *The Quarterly journal of economics* 135, 1399–1442.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2022b, Salience, *Annual Review of Economics* 14, 521–544.

Bose, Devdeepta, Henning Cordes, Sven Nolte, Judith Christiane Schneider, and Colin Farrell Camerer, 2022, Decision weights for experimental asset prices based on visual salience, *The Review of Financial Studies* 35, 5094–5126.

Brandstätter, Eduard, and Christof Körner, 2014, Attention in risky choice, *Acta psychologica* 152, 166–176.

Bruce, Neil DB, and John K Tsotsos, 2009, Saliency, attention, and visual search: An information theoretic approach, *Journal of vision* 9, 5–5.

Chen, Daniel L, Martin Schonger, and Chris Wickens, 2016, oTree-An open-source platform for laboratory, online, and field experiments, *Journal of Behavioral and Experimental Finance* 9, 88–97.

Chetty, Raj, Adam Looney, and Kory Kroft, 2009, Salience and taxation: Theory and evidence, *American economic review* 99, 1145–1177.

Cosemans, Mathijs, and Rik Frehen, 2021, Salience theory and stock prices: Empirical evidence, *Journal of Financial Economics* 140, 460–483.

Cubitt, Robin, Orestis Kopsacheilis, and Chris Starmer, 2022, An inquiry into the nature and causes of the description-experience gap, *Journal of Risk and Uncertainty* 65, 105–137.

d'Acremont, Mathieu, and Peter Bossaerts, 2016, Neural mechanisms behind identification of leptokurtic noise and adaptive behavioral response: Supplementary info 1, *Cerebral Cortex* 26, 1818–1830.

Danz, David, Lise Vesterlund, and Alistair J Wilson, 2022, Belief elicitation and behavioral incentive compatibility, *American Economic Review* 112, 2851–2883.

de Palma, André, Mohammed Abdellaoui, Giuseppe Attanasi, Moshe Ben-Akiva, Ido Erev, Helga Fehr-Duda, Dennis Fok, Craig R Fox, Ralph Hertwig, Nathalie Picard, et al., 2014, Beware of black swans: Taking stock of the description–experience gap in decision under uncertainty, *Marketing Letters* 25, 269–280.

Diecidue, Enrico, Haim Levy, and Moshe Levy, 2020, Probability dominance, *Review of Economics and Statistics* 102, 1006–1020.

Enke, Benjamin, and Thomas Graeber, 2023, Cognitive uncertainty, *The Quarterly Journal of Economics* 138, 2021–2067.

Enke, Benjamin, Thomas Graeber, and Ryan Oprea, 2023, Complexity and hyperbolic discounting .

Enke, Benjamin, Frederik Schwerter, and Florian Zimmermann, 2020, Associative memory and belief formation, Technical report, National Bureau of Economic Research.

Erev, Ido, and Greg Barron, 2005, On adaptation, maximization, and reinforcement learning among cognitive strategies., *Psychological review* 112, 912.

Erev, Ido, Eyal Ert, Alvin E Roth, Ernan Haruvy, Stefan M Herzog, Robin Hau, Ralph Hertwig, Terrence Stewart, Robert West, and Christian Lebiere, 2010, A choice prediction competition: Choices from experience and from description, *Journal of Behavioral Decision Making* 23, 15–47.

Erev, Ido, Ira Glozman, and Ralph Hertwig, 2008, What impacts the impact of rare events, *Journal of Risk and Uncertainty* 36, 153–177.

Fox, Craig R, and Liat Hadar, 2006, "Decisions from experience" = sampling error + prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004), *Judgement and Decision Making* 1, 159–161.

Frydman, Cary, and Milica Mormann, 2016, The role of salience and attention in choice under risk: An experimental investigation, *University of Southern California, Marshall School of Business* .

Graeber, Thomas, Florian Zimmermann, and Christopher Roth, 2022, Stories, statistics, and memory .

Gupta, Neeraja, Luca Rigotti, and Alistair Wilson, 2021, The Experimenters' Dilemma: Inferential Preferences over Populations.

Harel, Jonathan, Christof Koch, and Pietro Perona, 2006, Graph-based visual saliency, *Advances in neural information processing systems* 19.

Hartzmark, Samuel M, Samuel D Hirshman, and Alex Imas, 2021, Ownership, learning, and beliefs, *The Quarterly journal of economics* 136, 1665–1717.

Hertwig, Ralph, Greg Barron, Elke U Weber, and Ido Erev, 2004, Decisions From Experience and the Effect of Rare Events in Risky Choice, *Psychological Science* 15, 534–539.

Hertwig, Ralph, and Till Grüne-Yanoff, 2017, Nudging and boosting: Steering or empowering good decisions, *Perspectives on Psychological Science* 12, 973–986.

Hills, Thomas T, and Ralph Hertwig, 2010, Information search in decisions from experience: Do our patterns of sampling foreshadow our decisions?, *Psychological science* 21, 1787–1792.

Kahana, Michael Jacob, 2012, *Foundations of human memory* (OUP USA).

Kahneman, Daniel, and Amos Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47, 263–291.

Kaufmann, Christine, Martin Weber, and Emily Haisley, 2013, The Role of Experience Sampling and Graphical Displays on One's Investment Risk Appetite, *Management Science* 59, 323–340.

Khaw, Mel Win, Ziang Li, and Michael Woodford, 2021, Cognitive imprecision and small-stakes risk aversion, *The review of economic studies* 88, 1979–2013.

Laudenbach, Christine, Michael Ungeheuer, and Martin Weber, 2022, How to Alleviate Correlation Neglect in Investment Decisions, *Management Science* .

Li, Xiaomin, and Colin F Camerer, 2022, Predictable effects of visual salience in experimental decisions and games, *The Quarterly Journal of Economics* 137, 1849–1900.

Malmendier, Ulrike, and Stefan Nagel, 2011, Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?, *The Quarterly Journal of Economics* 126, 373–416.

Malmendier, Ulrike, and Jessica A Wachter, 2021, Memory of past experiences and economic decisions, *Available at SSRN 4013583* .

Markowitz, Harry M, 1952, Portfolio Selection, *Journal of Finance* 7, 77–91.

Mather, George, 2006, *Foundations of perception* (Taylor & Francis).

Miller, George A., 1956, The magical number seven, plus or minus two: Some limits on our capacity for processing information., *Psychological Review* 63, 81ÔÇô97.

Oprea, Ryan, 2022, Simplicity equivalents, *Working Paper* .

Paterson, Iain, and Andreas Diekmann, 1988, A Paradox in Decision Theory and Some Experimental Results: The Relative Nature of Decisions, *Theory and Decision* 25, 107–116.

Payzan-LeNestour, Elise, and Michael Woodford, 2022, Outlier blindness: A neurobiological foundation for neglect of financial risk, *Journal of Financial Economics* 143, 1316–1343.

Sharpe, William F., 1964, Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk, *Journal of Finance* 19, 425–442.

Spitmaan, Mehran, Emily Chu, and Alireza Soltani, 2019, Salience-driven value construction for adaptive choice under risk, *Journal of Neuroscience* 39, 5195–5209.

Taubinsky, Dmitry, and Alex Rees-Jones, 2018, Attention variation and welfare: theory and evidence from a tax salience experiment, *The Review of Economic Studies* 85, 2462–2496.

Tversky, Amos, and Daniel Kahneman, 1992, Advances in prospect theory: Cumulative representation of uncertainty, *Journal of Risk and Uncertainty* 5, 297–323.

Tversky, Amos, and Itamar Simonson, 1993, Context-dependent preferences, *Management science* 39, 1179–1189.

Ungeheuer, Michael, and Martin Weber, 2021, The Perception of Dependence, Investment Decisions, and Stock Prices, *Journal of Finance* 76, 797–844.

Ungeheuer, Michael, and Martin Weber, 2023, The Frequency of Outperformance and Investment Decisions.

Internet Appendix to

A Cognitive Foundation for Perceiving Uncertainty

Abstract

The Internet Appendix consists of five sections. We provide proofs and details about our theoretical model in Appendix A. Appendix B and C contain additional figures and tables, which are discussed in Appendix D. Complete instructions of our main experiments are reported in a supplementary document provided under this link.

A Theoretical Framework

This section presents a potential framework of bounded memory and attention and derives predictions on how these mechanisms interact with the learning environment to impact risky choice. We should note that this framework builds on several well-established aspects of limited attention and imperfect memory, such as the impact of salience on attribute weighting in judgment and choice as well as the impact of similarity and interference in memory recall. The model aims to formalize these aspects within a unified framework in order to explore the impact of the learning environment on decision-making, but this exercise should be viewed as illustrative given the potential alternative specifications that may capture the same psychological mechanisms, e.g., the salience function. We leave the more formal exploration of this topic to future work.

In the model, the DM first recalls information about an asset's distribution of outcomes and then evaluates this information when making a choice. Absent memory and attentional constraints, the DM's decision-problem reduces to the standard Expected Utility Theory model.

Consider assets A^i and states $s \in S$. Each state occurs with probability π_s. Asset A^i returns outcome x_s^i in state s. Let $v(x_s^i)$ denote the DM's subjective value from x_s^i. With no memory or attentional distortions, the DM will compute the following value $V^R(A^i)$ for each asset:

$$V^R(A^i) = \sum_{s \in S} \pi_s * v(x_s^i). \tag{4}$$

where the superscript R corresponds to the rational benchmark. A DM in our setting will depart from this baseline due to distortions in the perception of the distributions due to a) imperfect recall of outcomes because of memory constraints and b) over/underweighting of the likelihoods because of attentional constraints. To preview the results, recall will be adversely affected by the number of states as well as the similarity of outcomes across states and assets; over/underweighting will be affected by the salience of states, which is a function of both the payout distribution and visual factors.

A.1 Gathering information

We consider the choice between two assets, A^U and A^F. To make a choice, the DM assembles information about the distribution of risk either from being presented with it simultaneously or by learning sequentially. In the case of the latter, she must accrue this information through memory recall. The process of memory recall in our framework builds on the models of Bordalo et al. (2020) and Anderson et al. (1998). Let $S(u, v) : E \times E \to [0, 1]$ measure the similarity between two outcomes u and v in a DM's set of sequentially sampled outcomes E, where $S(u, v) = 1 \iff u = v$. Let χ_s^i be the remembered outcome in state s for asset A^i. The DM correctly recalls a given outcome (i.e., $\chi_s^i = x_s^i$) with probability

$$r(x_s^i) = \sqrt{\frac{S(x_s^i, x_s^i)}{\sum_k S(x_s^i, x_s^k)} * \frac{S(x_s^i, x_s^i)}{\sum_t S(x_s^i, x_t^i)}} \tag{5}$$

$$= \sqrt{\frac{1}{\sum_k S(x_s^i, x_s^k)} * \frac{1}{\sum_t S(x_s^i, x_t^i)}}. \tag{6}$$

With probability $1 - r(x_s^i)$, the DM instead recalls the mean of the two outcomes most similar to the recall target in that state. That is, the DM experiences interference in her ability to recall outcomes from two sources: the similarity of outcomes across assets within state and the similarity of outcomes across states within asset. Recall is perfect when outcomes are entirely dissimilar across both states and assets, and recall deteriorates as either type of similarity increases.

A.2 Forming mental representation

Once the DM has access to the distribution, either through the recall process described above or by being presented with it directly, she forms a mental representation of the outcomes and respective likelihoods. This mental representation is subject to attentional distortions as a function of state-specific salience. Namely, states that are more (less) salient will attract more (less) attention and be over(under)weighted in the decision-making process.

A large literature in economics and psychology has shown the importance of attention in risky choice. Limited attention is channeled towards some states over others as a function of their salience (Brandstätter and Körner, 2014; Spitmaan et al., 2019; Bordalo et al., 2022b). As a result of channeled attention, states that are more salient are overweighted in the decision-making process

compared to states that are less salient (Frydman and Mormann, 2016). In the context of risky choice, a state's salience is determined by both visual stimuli (Frydman and Mormann, 2016; Bose et al., 2022) and payoffs (Bordalo et al., 2013).[22]

When forming a mental representation, the weight assigned to a state is a function of its salience. Specifically, in the framework, states are first ranked by their relative salience. Then, objective state probabilities are replaced with distorted probabilities, where more salient states are overweighted and less salient states are underweighted. Let $\sigma(\chi_s^i, \chi_s^j, O_s)$ be a function mapping recalled outcomes to state salience and k_s be the relative ranking of state s by salience. We take the salience to be a function of 1) relative payoffs (x_s^i, x_s^j) and 2) visual stimuli associated with the state O_s.

Following Bruce and Tsotsos (2009), the salience of a state is a function of a bottom-up, stimulus-driven process that is part of sampling behavior. In their Attention based on Information Maximization (AIM) model, salience emerges from first principles of bounded rationality, namely efficient coding and information theory. The salience of some component (in our case, the state) is a function of the local information relative to the surrounding information; a state is more salient if the content is unexpected given the surrounding context.

For risky choice, an important factor that determines the degree to which a state departs from a local context is its associated outcomes. We follow Bordalo et al. (2012) in capturing salience through outcomes. The salience of state s for asset A^i is a continuous and bounded function $\sigma(x_s^i, x_s^j)$ that satisfies:

1. Ordering. If for states $s, \tilde{s} \in S$ we have that $[x_s^{\min}, x_s^{\max}] \subset [x_{\tilde{s}}^{\min}, x_{\tilde{s}}^{\max}]$ then

$$\sigma(x_s^i, x_s^j, O_s) < \sigma(x_{\tilde{s}}^i, x_{\tilde{s}}^j, O_{\tilde{s}}).$$

2. Diminishing sensitivity. If $x_s^j > 0$ for $j = 1, 2$, then for any $\epsilon > 0$,

$$\sigma(x_s^i + \epsilon, x_s^j + \epsilon, O_s) < \sigma(x_s^i, x_s^j, O_s).$$

[22] For empirical evidence on the importance of state-based salience in the domain of belief-updating, see Ba et al. (2022).

3. Reflection. For any two states $s, \tilde{s} \in S$ such that $x_s^j, x_{\tilde{s}}^j > 0$ for $j = 1, 2$, we have

$$\sigma(x_s^i, x_s^j, O_s) < \sigma(x_{\tilde{s}}^i, x_{\tilde{s}}^j, O_{\tilde{s}}) \Leftrightarrow \sigma(-x_s^i, -x_s^j, O_s) < \sigma(-x_{\tilde{s}}^i, -x_{\tilde{s}}^j, O_{\tilde{s}}).$$

As shown in Li and Camerer (2022), visual stimuli are also an important aspect of salience. A state whose outcomes depart strongly from the local payoffs of other states—either due to contrasting colors or pattern-based gestalt (Mather, 2006)—is more likely to be salient than a state that departs less from the local context. Let O_s correspond to the visual salience parameter such that $O_s \in [0, 1]$ and $\sum_{s \in S} O_s = 1$. Conditional on the same outcomes $(x_s^i, x_s^j) = (x_{\tilde{s}}^i, x_{\tilde{s}}^j)$, a state s is ranked higher in salience than $\tilde{s}$ if its visual salience is higher $O_s > O_{\tilde{s}}$ based on the local context.

Notably, as shown in Bose et al. (2022), Li and Camerer (2022), Bruce and Tsotsos (2009), amongst others, the salience of an object can be predicted using algorithms such as the Graph-Based Visual Saliency model (Harel et al., 2006). This technique can be used to separate which aspects of saliency are driven purely by the payoffs versus other types of visual stimuli (e.g., colors, relative placement in a field).

After ranking the states based on salience, the weight then given to state s is:

$$\omega_s = \frac{\delta^{k_s}}{\sum_{r \in S} \delta^{k_r} * \pi_r}, \tag{7}$$

where $\delta \in (0, 1]$. Finally, the DM evaluates each asset by computing:

$$V(A^i) = \sum_{s \in S} \hat{\pi}_s^i * v(\chi_s^i), \tag{8}$$

where $\hat{\pi}_s^i = \pi_s * \omega_s$.

A.3 Choice behavior

Consider a decision environment with a choice between two assets A^U and A^F and a finite number of equally-likely states $|S|$, where $x_s^F > x_s^U$ for all but one state s^*, where $x_s^F < x_s^U$. Finally, as in our experimental setting, for each outcome in A^U, there exists an equivalent outcome in some state for A^F and vice-versa. A key property of this environment is that $V^R(A^U) = V^R(A^F)$. Our first prediction is on the choice between assets under perfect recall.

Prediction 1. *With no memory constraints, a DM with salience-distorted attention will prefer A^U to A^F.*

Proof. We start by showing that the salience σ of state s^* is greater than the salience of any other state s_j. Note that since the DM is assumed to be affected by salience, $\delta < 1$.

$$\forall s \in S, s \neq s^*, x_s^U > x_s^F \tag{9}$$

$$x_*^U > x_*^F \tag{10}$$

$$\forall s \in S,\ \exists t \text{ s.t. } x_s^i = x_t^j \tag{11}$$

$$\implies \sum_s x_s^U = \sum_s x_s^F \tag{12}$$

$$\implies \forall s \neq s^* [x_*^F, x_*^U] \supset [x_s^U, x_s^F] \tag{13}$$

$$\implies \sigma(x_*^i, x_*^j) > \sigma(x_s^i, x_s^j). \tag{14}$$

Thus, we know the relative order of salience with respect to state, which allows us to compare ω between each state:

$$k_* = 1 \tag{15}$$

$$\forall s \neq s^*, k_s > 1 \tag{16}$$

$$\implies \delta^{k_*} > \delta^{k_s} \tag{17}$$

$$\implies \frac{\delta^{k_*}}{\sum_{r \in S} \delta^{k_* \pi_r}} > \frac{\delta^{k_s}}{\sum_{r \in S} \delta^{k_s \pi_r}} \tag{18}$$

$$\implies \omega_* > \omega_s. \tag{19}$$

Since the DM has no memory constraints, $\chi_s^i = x_s^i$. Recalling our valuation function,

$$V(A^i) = \sum_{s \in S} \hat{\pi}_s^i * v(\chi_s^i) \tag{20}$$

$$= \sum_{s \in S} \hat{\pi}_s^i * v(x_s^i). \tag{21}$$

Since each state is equally likely,

$$V(A^i) = \sum_{s \in S} \omega_s * v(x_s^i). \tag{22}$$

Recalling equations (5) and (6), we see that

$$\sum_{s \in S} v(x_s^U) = \sum_{s \in S} v(x_s^F). \tag{23}$$

Here, we assume that $v(\cdot)$ is monotone increasing in its domain. Since $\omega_* > \omega_s$ and $x_*^F < x_*^U$,

$$\sum_{s \in S} \omega_s * v(x_s^U) > \sum_{s \in S} \omega_s * v(x_s^F) \tag{24}$$

$$\implies V(A^U) > V(A^F) \tag{25}$$

□

This prediction is directly relevant to the "simultaneous" learning environment, where the DM is presented with information about the risk distributions directly. There, imperfect recall does not play a role since all of the information necessary to form an accurate mental representation is present at the time of choice.

Our second prediction introduces the possibility of imperfect recall of the relevant outcomes.

Prediction 2. *Consider a setting with imperfect recall. Suppose the similarity between outcomes is nonzero across both states and assets. Then there exists a number of states $|S^*|$, such that for $|S| > |S^*|$, the DM will prefer A^F to A^U.*

Proof. Let $r(x_s^i) \in (0, 1)$ be the probability of correctly recalling the outcome of asset A^i in state s. We start by highlighting that the preference for A^U in the unconstrained-memory environment is driven by the fact that $x_*^U > x_*^F$. Let $0 < \zeta < 1$ be the similarity between asset outcomes for A^U and A^F within state and $0 < \underline{\psi_s} \leq \overline{\psi_s} < 1$ be the minimum and maximum similarity, respectively, between asset outcomes within asset across states. We can then establish upper and lower bounds

for $r(x_s^i)$:

$$r(x_s^i) = \sqrt{\frac{1}{\sum_k S(x_s^i, x_s^k)} * \frac{1}{\sum_t S(x_s^i, x_t^i)}} \tag{26}$$

$$\implies 0 < \sqrt{\frac{1}{1+\zeta} * \frac{1}{1+\overline{\psi_s}(|S|-1)}} \leq r(x_s^i) \leq \sqrt{\frac{1}{1+\zeta} * \frac{1}{1+\underline{\psi_s}(|S|-1)}} < 1. \tag{27}$$

Now consider the decision environment. As shown in Bordalo et al. (2012), a DM unaffected by memory prefers A^U to A^F if and only if

$$\sum_{s \in S} \delta^{k_s} \pi_s [v(x_s^U) - v(x_s^F)] > 0. \tag{28}$$

In our setting, this becomes

$$\sum_{s \in S} \delta^{k_s} [v(\chi_s^U) - v(\chi_s^F)] > 0. \tag{29}$$

Prior to any memory distortions, $v(x_s^U) < v(x_s^F)$ for all but s^*, where $v(x_*^U) > v(x_*^F)$. For any given state, each outcome can either be remembered correctly with probability ρ, or be replaced by the average of the two outcomes in that state with probability $(1-\rho)$. Thus, one possible outcome is that the previously dominant asset remains dominant. This is the case if just one or neither state is misremembered, which happens with probability $(2\rho - \rho^2)$. The other possibility is that both outcomes are misremembered, in which case the state falls out of consideration; this happens with probability $(1-\rho)^2$.

Turning our attention to the cumulative realization of the recall process, we consider three main possibilities:

1. The DM remembers either or both outcomes from the originally salient state s^*, in which case her preference is ambiguous. This occurs with probability $p_A = (2r - r^2)$.

2. The DM forgets both outcomes from the salient state s^* but remembers at least one other outcome in some state, in which case she prefers asset A^F to asset A^U. This occurs with probability $p_F = (1-r)^2 - (1-r)^{2|S|}$.

3. The DM forgets every outcome from every state, in which case she is indifferent. This occurs

with probability $p_N = (1-r)^{2|S|}$.

Rewriting these probabilities using the inequalities from (29), we get:

$$p_A \leq 2\sqrt{\frac{1}{1+\zeta} * \frac{1}{1+\underline{\psi_s}(|S|-1)}} - \left(\frac{1}{1+\zeta} * \frac{1}{1+\overline{\psi_s}(|S|-1)}\right) \tag{30}$$

$$p_F \geq \left(1 - \sqrt{\frac{1}{1+\zeta} * \frac{1}{1+\underline{\psi_s}(|S|-1)}}\right)^2 - \left(1 - \sqrt{\frac{1}{1+\zeta} * \frac{1}{1+\overline{\psi_s}(|S|-1)}}\right)^{2|S|} \tag{31}$$

$$p_N \leq \left(1 - \sqrt{\frac{1}{1+\zeta} * \frac{1}{1+\overline{\psi_s}(|S|-1)}}\right)^{2|S|} \tag{32}$$

Finally, note that $p_F \to 1$ as $|S| \to \infty$ while both p_A and p_N approach zero. Therefore, we can find some $|S|$ such that $p_F > 0.5$.

□

The intuition for the second prediction is straightforward. If the DM misremembers the states in which $x_s^U > x_s^F$, while remembering the states where $x_s^U < x_s^F$, then she will prefer A^F to A^U. Consider the structure of A^U and A^F, where the former outperforms the latter in one state while the latter outperforms the former in the others. This implies that as the number of states grows, the chance that the DM will misremember the state in which $x_s^U > x_s^F$ while remembering a state where $x_s^U < x_s^F$ will increase as well. Importantly, because imperfect recall plays a bigger role in the sequential information environment than in the simultaneous information environment, the first two predictions imply a preference reversal from A^F being preferred to A^U in the former to A^F being preferred to A^U in the latter.

Our third prediction is on how similarity impacts choice behavior through the recall process.

Prediction 3. *Consider a setting with imperfect recall and a total number of states $|S|$. The likelihood that A^F is preferred to A^U is increasing in the similarity of outcomes both across assets within a state and across states within an asset.*

Proof. Since we have shown in the above proof that p_F is increasing in $1/r$, this prediction follows from the fact that either type of similarity reduces recall probability. □

The intuition for the third prediction follows from the second. Similarity between outcomes increases the chance that they will be misremembered. The relative impact of similarity is higher for infrequent states; namely, while higher similarity will decrease recall accuracy across all states, it decreases the chances of misremembering a state where $x_s^U > x_s^F$ than the chances of misremembering a state $x_s^U < x_s^F$, simply because there are more of the latter. Thus, conditional on a number of states $|S|$, the DM is (weakly) more likely to prefer A^F to A^U as the similarity between outcomes increases.

We build on this intuition for our last prediction, which directly tests the interference aspect of imperfect recall. Consider a third asset A^D which the DM learns about but cannot choose. Interference predicts that one could reverse the preference between assets in the sequential information environment when the third asset is more similar to asset A^U than asset A^F.

Prediction 4. *Suppose a third asset A^D is included, which dominates both A^D and A^F in every state, i.e., $x_s^D = x_s^{max} + \epsilon_s$, $\epsilon_s > 0 \ \forall s \in S$. The likelihood that A^U is preferred to A^F is increasing in the similarity between the outcomes of A^D and A^U.*

Proof. We first note that, although a new asset has been added, the ranking of each state by salience remains the same. This prediction follows from the fact that, in this case, forgetting any x_s^U increases $v(\chi_s^U) - v(\chi_s^F)$, so

$$\mathbb{E}\left[\sum_{s\in S} \delta^{k_s}[v(\tilde{\chi}_s^U) - v(\tilde{\chi}_s^F)]\right] > \mathbb{E}\left[\sum_{s\in S} \delta^{k_s}[v(\chi_s^U) - v(\chi_s^F)]\right],$$

where $S(\tilde{x}_s^U, \tilde{x}_s^D) > S(x_s^U, x_s^D)$. □

The logic for this prediction comes from the property of interference that underlies most models of memory. When an outcome from a state is misremembered, the outcome that replaces it is more likely to be associated with a similar source. If the "replacement" outcomes from the similar source are superior to those of both assets in the choice set, then misremembering an outcome from A^U will increase its attractiveness. Including the similar asset also increases the chances that an outcome from A^U is misremembered in the first place. Both forces lead to an increased preference for A^U in the sequential information learning environment.

For our last prediction, consider A^D that is more similar to A^F than A^U.

Prediction 5. *Suppose a third asset A^D is included, where $x_s^D = x_s^F + \epsilon_s$, $\epsilon_s > 0 \; \forall s \in S$. When choosing between A^U and A^F, the likelihood that A^F is preferred to A^U is increasing in the similarity of x_s^D to x_s^F.*

Proof. We first note that, again, although a new asset has been added, the ranking of each state by salience remains the same. The prediction then follows from the fact that forgetting any x_s^F increases $v(\chi_s^F) - v(\chi_s^U)$, so

$$\mathbb{E}\left[\sum_{s\in S} \delta^{k_s}[v(\tilde{\chi}_s^F) - v(\tilde{\chi}_s^U)]\right] > \mathbb{E}\left[\sum_{s\in S} \delta^{k_s}[v(\chi_s^F) - v(\chi_s^U)]\right],$$

where $S(\tilde{x}_s^F, \tilde{x}_s^D) > S(x_s^F, x_s^D)$. □

Note that the condition that asset A^F is preferred to A^U when the third asset A^D is more similar to the former only requires A^D to dominate A^F, not both A^F and A^U.

B Additional Figures

Figure IA1: Flow of the experiments. This figure illustrates the flow of our Baseline Experiment. It begins with instructions, which include information on the upcoming tasks (choices and belief elicitation), information to come (presentation of risky alternatives), as well as an explanation of how participants are paid (performance-based incentivization either for their choice or stated beliefs). Subsequently participants go through two rounds, where they choose between two alternatives and state beliefs. Specifically, they are first shown information on the outcomes of two risky alternatives (or 'assets'). The learning environment is randomly allocated between participants, including both simultaneous information environments (stating outcomes for all possible states and the probabilities) and sequential information environments (where participants draw a representative sample from the underlying joint distribution). Then, participants make a choice between the assets and/or state their beliefs about the expected outcomes and likelihoods of outcomes. Participants make their choice on a first screen and state their beliefs on a second screen if they are incentivized for their choice. If they are incentivized for accurate beliefs, they state their beliefs on a first screen and make their choice on the second screen. After finishing the second round, participants are asked a few concluding questions, before they are informed about their payment.

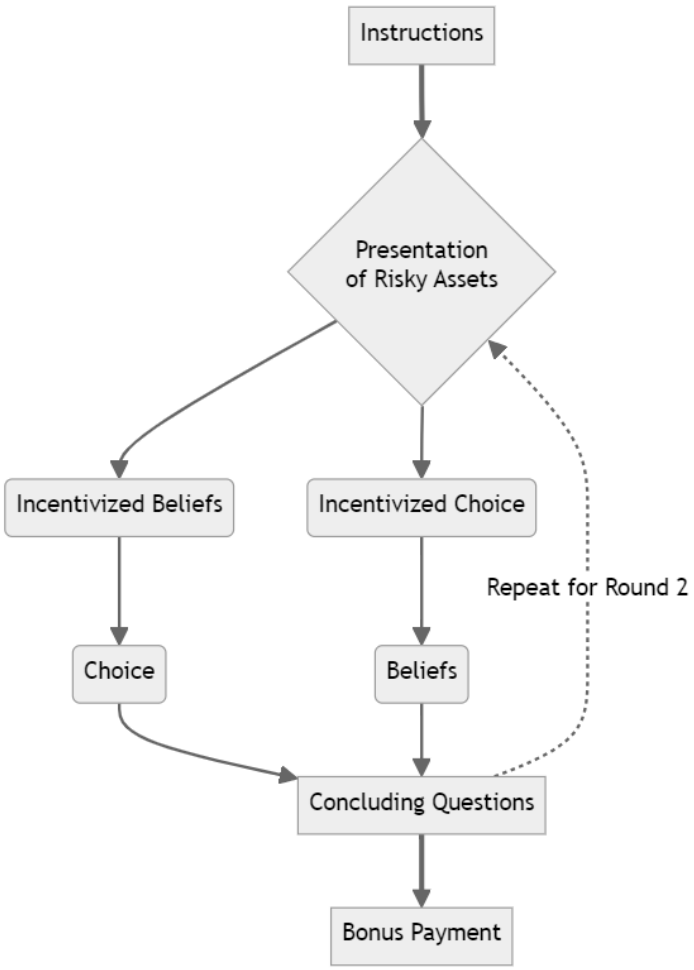


Figure IA2: Baseline Experiment: Simultaneous Information Environment for the 51-State Round. This figure shows the simultaneous information graph, which is used in our Baseline Experiment to present participants in the simultaneous information learning environment with the two risky assets when they make choices / state beliefs (see 'Presentation of Risky Assets' in Figure IA1). The graphical and verbal description are analogous in the 11-state round. This figure is shown shown in the simultaneous information environment on the choice / beliefs screens.

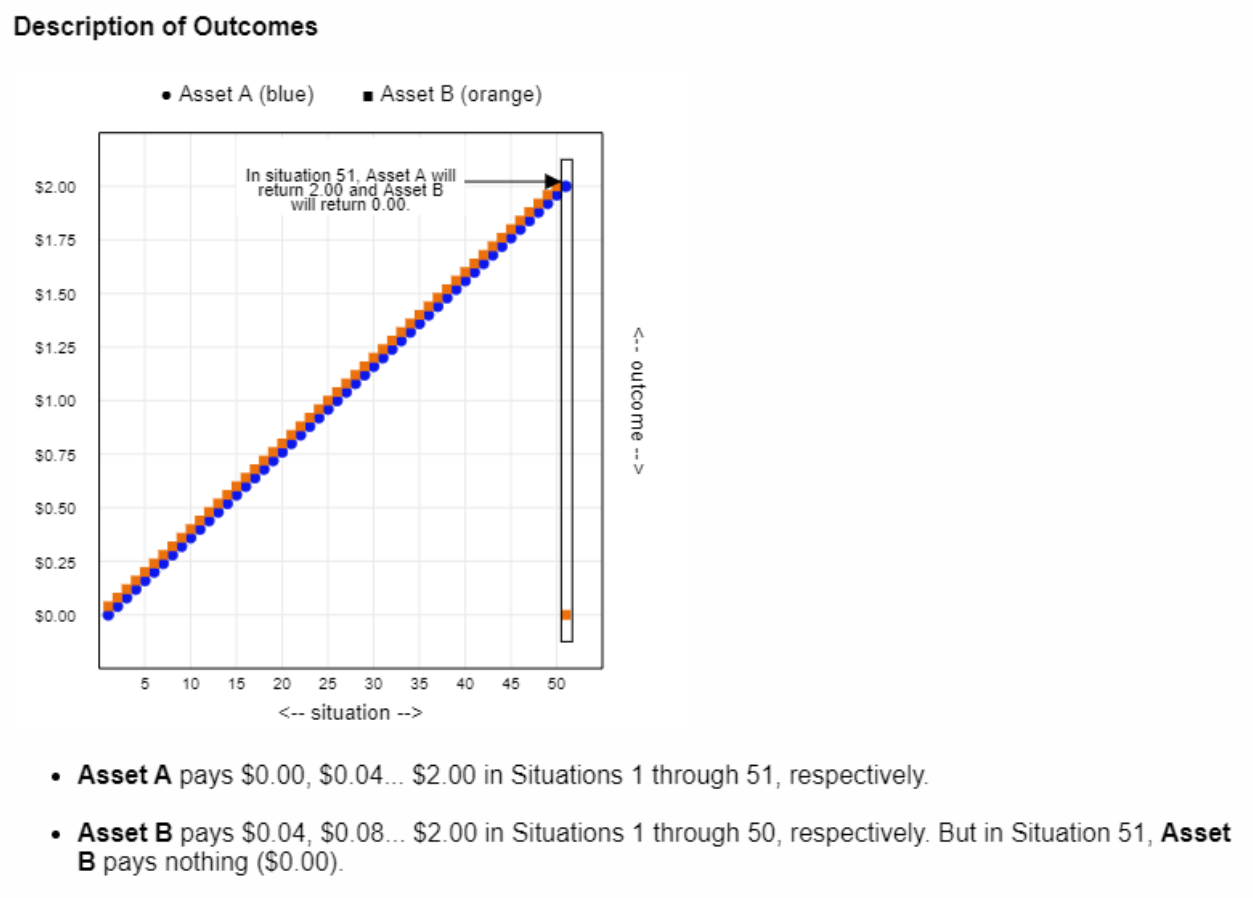


Figure IA3: Sequential Information with "No Recall Bias" – Change from Baseline. This figure illustrates the changes we made for the sequential information treatment in the "No Recall Bias" condition.

Panel A: Early Announcement Screen

You will be asked to write down and keep track of some of the information that you see in this experiment. Please get a pen and paper or other method of writing down the information now.

Panel B: Detailed Explanation Before Sequential Information

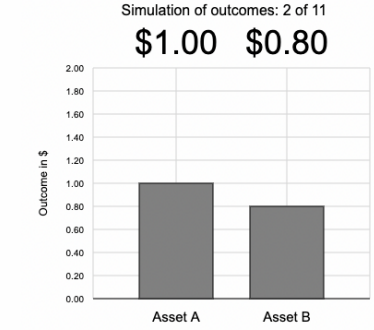
We will now begin the experiment. In this round, there are two Assets A and B and 11 possible situations, which are equally likely. On the screens that follow, you will see a simulation of how each Asset does in each of the 11 possible situations. This will give you an impression of how the assets perform. The outcomes are drawn from the joint distribution of outcomes across the two assets.

The bonus will be based on the accuracy of one of your guesses. The guess for your payment will be randomly selected.

The graph below depicts a typical simulation screen. In this situation, Asset A generated an outcome of 1.00 while Asset B generated an outcome of 0.80. You will see a graph like this across all 11 possible situations.

Please write down the outcomes for every single situation that you are shown, 11 total. For example, in the situation below, you can write down "Situation 2, Asset A returned $1.00 and Asset B returned $0.80." You will be asked to recall this information at the end of the study and be paid a bonus of $1 if you recall this information correctly.

You will then make decisions based on this information.



Panel C: Opportunity for Bonus After Choice & Beliefs

For each situation, please report below which asset returned more in that situation (or if both assets returned the same amount). If you answer each question correctly, you will receive a bonus of $1.

	Asset A returned more than Asset B.	Asset B returned more than Asset A.	Both assets returned the same amount.
Situation 1	○	○	○
Situation 2	○	○	○
Situation 3	○	○	○
Situation 4	○	○	○
Situation 5	○	○	○
Situation 6	○	○	○
Situation 7	○	○	○
Situation 8	○	○	○
Situation 9	○	○	○
Situation 10	○	○	○
Situation 11	○	○	○

Figure IA4: Baseline Experiment: Viewing Time Split. This figure shows the fraction of participants who select each asset after the sequential information learning environment, split by whether their viewing time for the extreme outcome pair was higher or lower than their average viewing time for the moderate outcome pairs. The result for the 71% (29%) of participants with a viewing time for the extreme pair above (below) their average viewing time for the moderate pairs is shown on the right (left). A significantly higher fraction of participants selects the frequently-outperforming alternative than the unlikely-outperforming alternative, even for the majority of participants who pay over-proportional attention to the extreme outcome pair.

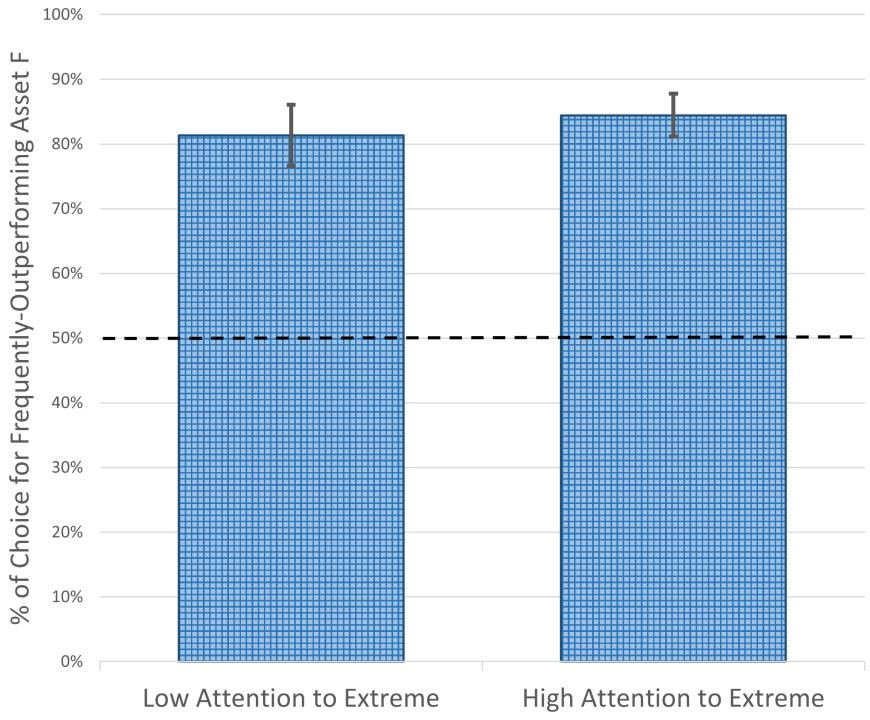
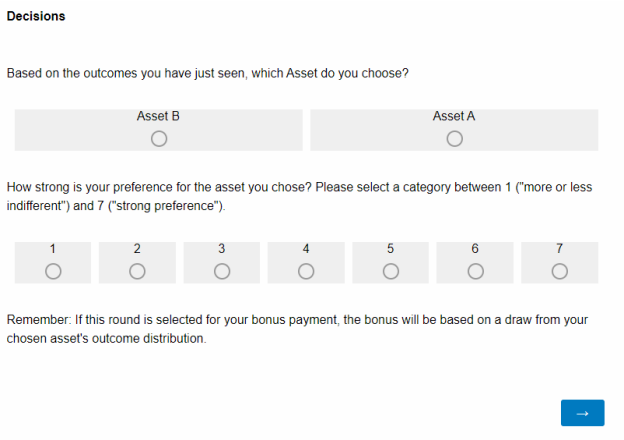


Figure IA5: Baseline Experiment: Elicitation of Choices and Beliefs in the Choice-Sequential Information Group. This figure shows the screens, which are used to ask participants who see the sequential information environment for their choice and beliefs in the choice group (first asked and incentivized for their choice, see Panel A, then for beliefs, see Panel B).

Panel A: Choice screen



Panel B: Beliefs screen

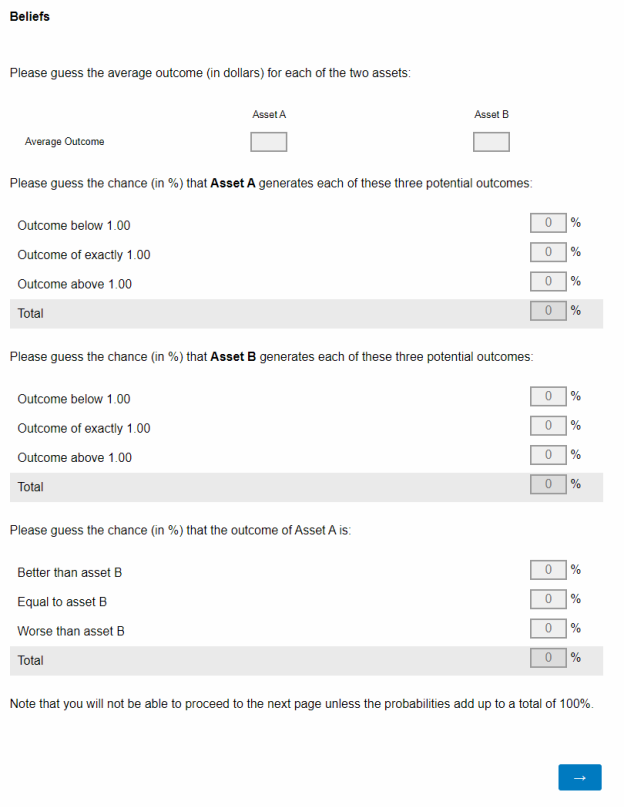
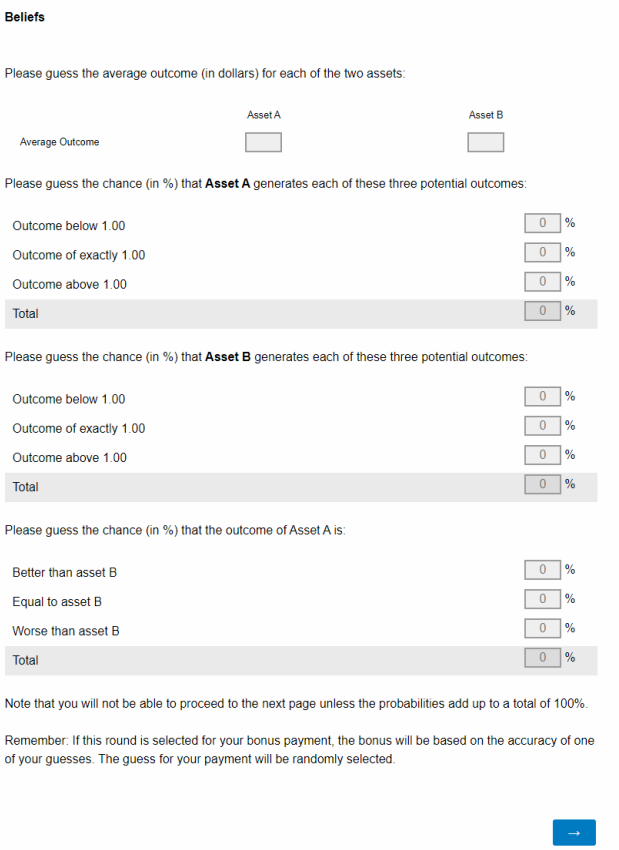


Figure IA6: Baseline Experiment: Elicitation of Beliefs and Choices in the Beliefs-Sequential Information Group. This figure shows the screens, which are used to ask participants who see the sequential information environment for their beliefs and choice in the beliefs group (first asked and incentivized for their beliefs, see Panel A, then for choice, see Panel B).

Panel A: Beliefs screen



Panel B: Choice screen

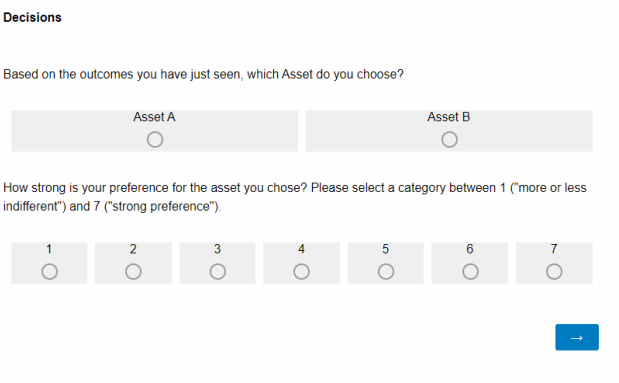
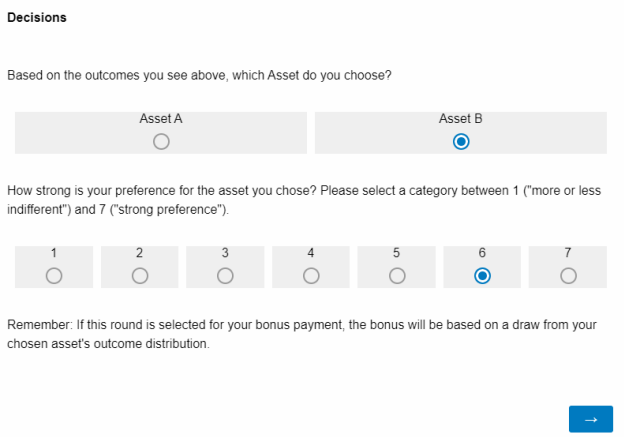


Figure IA7: Baseline Experiment: Elicitation of Choices and Beliefs in the Choice-Simultaneous Information group. This figure shows the screens, which are used to ask participants who see the simultaneous information environment for their choices and beliefs in the choice group (first asked and incentivized for their choice, see Panel A, then for beliefs, see Panel B).

Panel A: Choice screen



Panel B: Beliefs screen

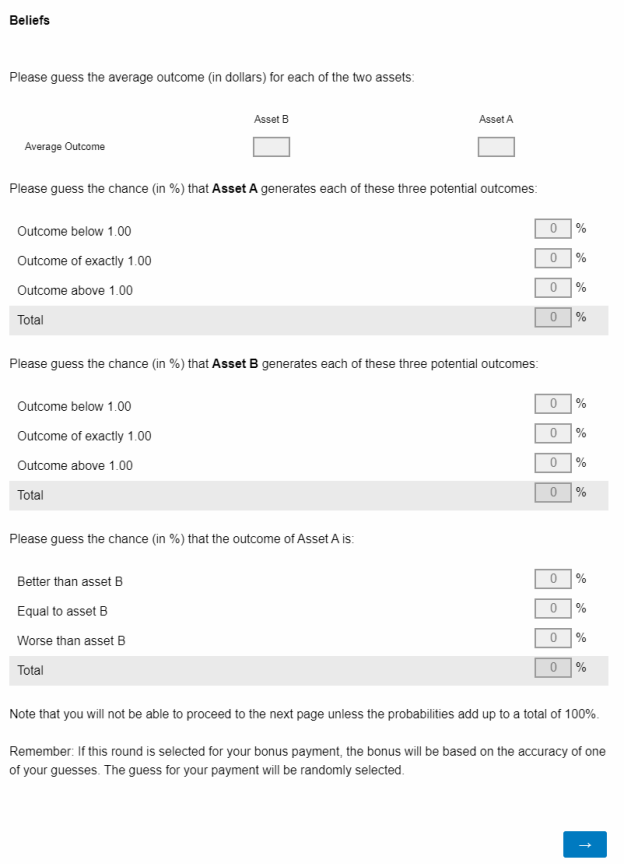
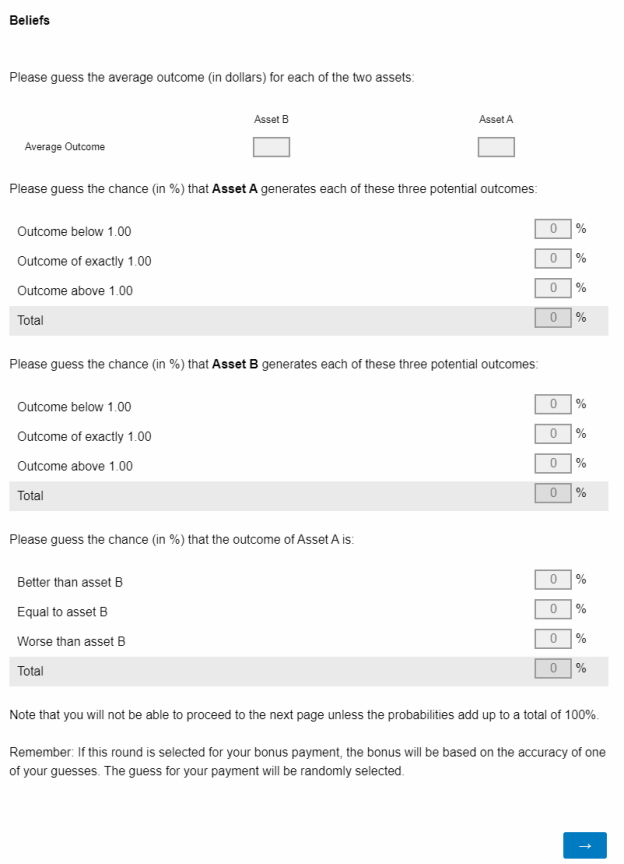


Figure IA8: Baseline Experiment: Elicitation of Beliefs and Choices in the Beliefs-Simultaneous Information group. This figure shows the screens, which are used to ask participants who see the simultaneous information environment for their beliefs and choice in the beliefs group (first asked and incentivized for their beliefs, see Panel A, then for choice, see Panel B).

Panel A: Beliefs screen



Panel B: Choice screen

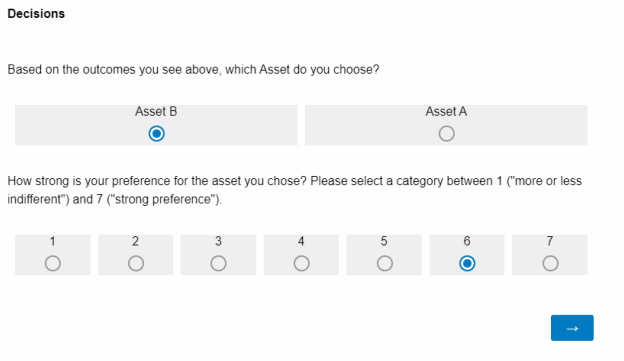


Figure IA9: Baseline Experiment: Recency or Primacy effects. This figure shows the fraction of participants who select the frequently-outperforming asset after sequential information, split by when the extreme outcome pair was sampled. The split is done by quintile, from a bin with the first 20% to a bin with the last 20% of the sequence. The middle quintile includes one more observation (e.g., observations 5-7 in the short 11-state round, where other quintiles are 1-2, 3-4, 8-9, and 10-11). Point estimates are fractions of participants that select the frequently-outperforming (but infrequently much worse) alternative. 95% confidence intervals are displayed for each point estimate, showing that all point estimates are significantly higher than 50%.

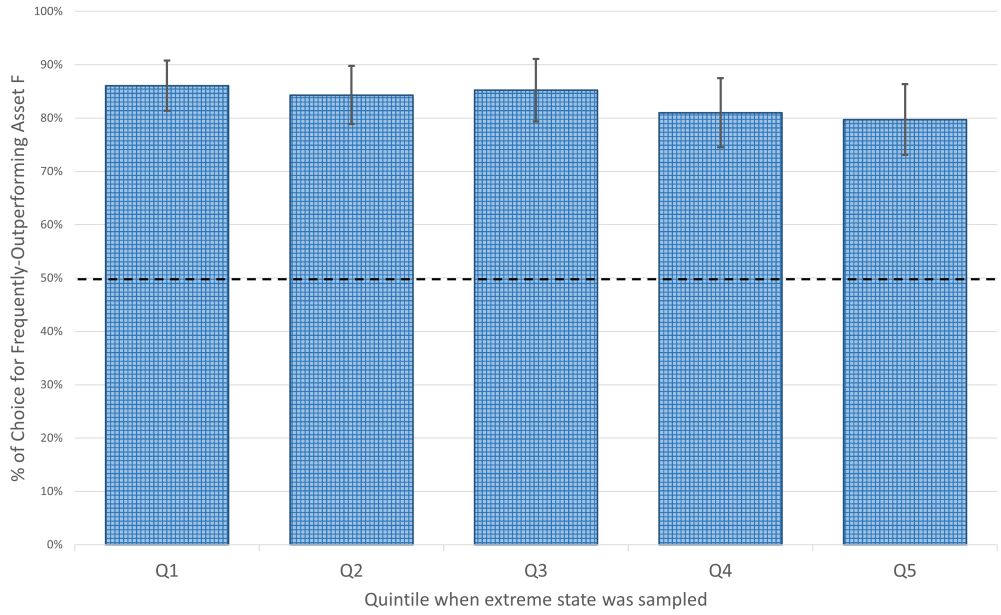


Figure IA10: Experiment A1: Verbal Simultaneous Information Environment. This figure illustrates the learning environment in the verbal simultaneous information group of Experiment A1, which is used to present participants with the two assets in the 'Information' stage (see dotted box in Figure IA1). The order of bullet points for Alternative 2 is randomized, so that the extreme outcome pair is randomly at the top or bottom. The allocation of the frequently-outperforming asset to Alternative 1 or 2 is also randomized.

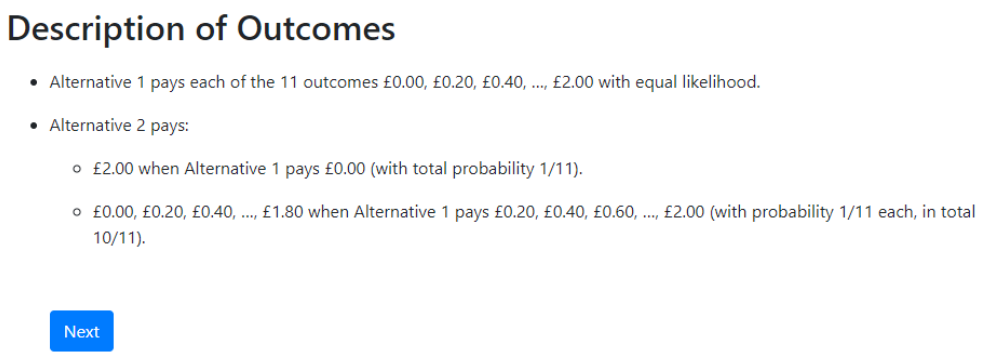
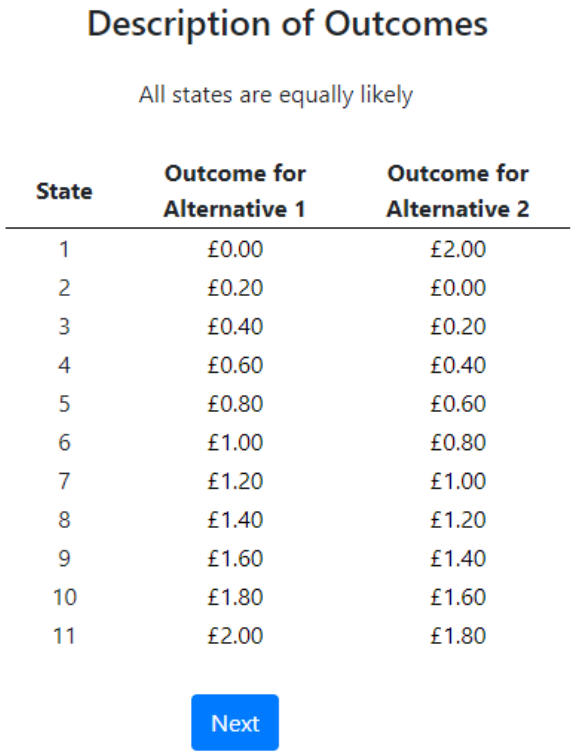


Figure IA11: Experiment A1: Tabular Simultaneous Information Environment. This figure illustrates the learning environment in the tabular simultaneous information group of Experiment A1, which is used to present participants with the two assets in the 'Information' stage (see dotted box in Figure IA1). Outcome pairs are sorted either by Alternative 1 or 2, so that the extreme outcome pair is randomly at the top or bottom. The allocation of the frequently-outperforming asset to Alternative 1 or 2 is also randomized.



Description of Outcomes

All states are equally likely

State	Outcome for Alternative 1	Outcome for Alternative 2
1	£0.00	£2.00
2	£0.20	£0.00
3	£0.40	£0.20
4	£0.60	£0.40
5	£0.80	£0.60
6	£1.00	£0.80
7	£1.20	£1.00
8	£1.40	£1.20
9	£1.60	£1.40
10	£1.80	£1.60
11	£2.00	£1.80

Next

Figure IA12: Experiment A1: Verbal Simultaneous Environment with Differences. This figure illustrates the learning environment in the verbal description group with differences of Experiment A1, which is used to present participants with the two assets in the 'Information' stage (see dotted box in Figure IA1). The order of bullet points for Alternative 2 is randomized, so that the extreme outcome pair is randomly at the top or bottom. The allocation of the frequently-outperforming asset to Alternative 1 or 2 is also randomized.

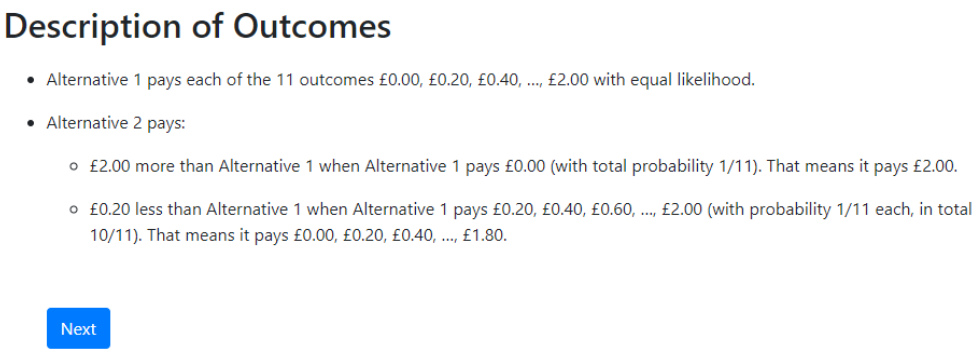
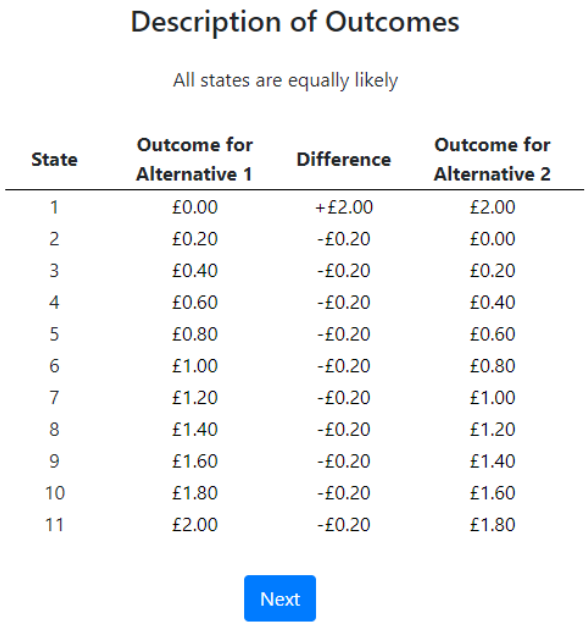


Figure IA13: Experiment A1: Tabular Simultaneous Information Environment with Differences. This figure illustrates the learning environment in the tabular simultaneous information with differences group of Experiment A1, which is used to present participants with the two assets in the 'Information' stage (see dotted box in Figure IA1). Outcome pairs are sorted either by Alternative 1 or 2, so that the extreme outcome pair is randomly at the top or bottom. The allocation of the frequently-outperforming asset to Alternative 1 or 2 is also randomized.



Description of Outcomes

All states are equally likely

State	Outcome for Alternative 1	Difference	Outcome for Alternative 2
1	£0.00	+£2.00	£2.00
2	£0.20	-£0.20	£0.00
3	£0.40	-£0.20	£0.20
4	£0.60	-£0.20	£0.40
5	£0.80	-£0.20	£0.60
6	£1.00	-£0.20	£0.80
7	£1.20	-£0.20	£1.00
8	£1.40	-£0.20	£1.20
9	£1.60	-£0.20	£1.40
10	£1.80	-£0.20	£1.60
11	£2.00	-£0.20	£1.80

Next

Figure IA14: Experiment A1: Sequential Information Environment with Differences. This figure illustrates the learning environment in the sequential information group with differences of Experiment A1, which is used to present participants with the two assets in the 'Information' stage (see dotted box in Figure IA1). Subjects sample T outcome pairs for the two assets (T=11 or 101 in Experiment A1's two rounds, counterbalanced), before moving on to the tasks of the experiment (choice and belief elicitation). Outcome pairs make up a representative distribution and are sampled (without replacement) in random order for each participant. The allocation of the frequently-outperforming asset to Alternative 1 or 2 is also randomized.

Panel A: Illustration sequential information screen, moderate outcome pair



Panel B: Illustration summary screen, extreme outcome pair

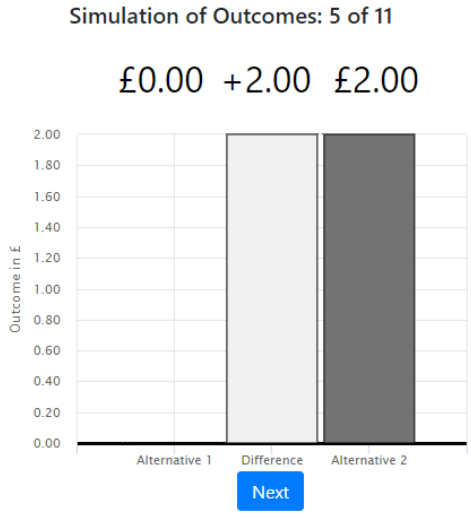


Figure IA15: Experiment A1: Sequential Information Environment. This figure illustrates the learning environment in the sequential information group of Experiment A1, which is used to present participants with the two assets in the 'Information' stage (see dotted box in Figure IA1). Subjects sample T outcome pairs for the two assets (T=11 or 101 in Experiment A1's two rounds, counterbalanced), before moving on to the tasks of the experiment (choice and belief elicitation). Outcome pairs make up a representative distribution and are sampled (without replacement) in random order for each participant. The allocation of the frequently-outperforming asset to Alternative 1 or 2 is also randomized.

Panel A: Illustration sequential information screen, moderate outcome pair

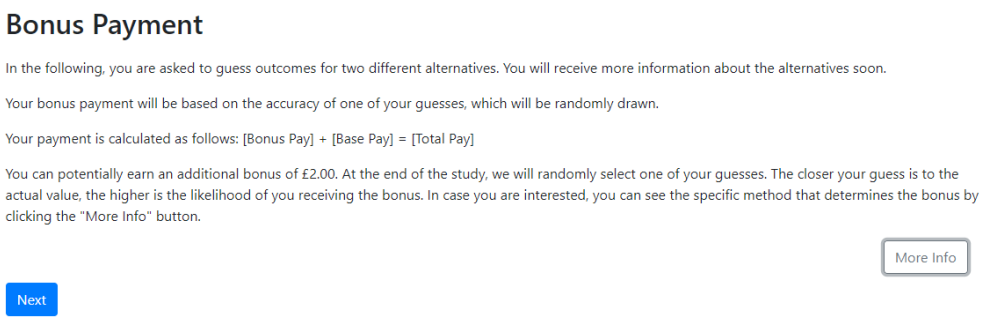


Panel B: Illustration summary screen, extreme outcome pair



Figure IA16: Experiment A2: Instructions Beliefs Group. This figure illustrates the instructions received by participants in the beliefs group of Experiment A2 before the start of the first round. The pop-up window in Panel B appears when participants click on 'More info'.

Panel A: Main instructions



Panel B: Pop-up window

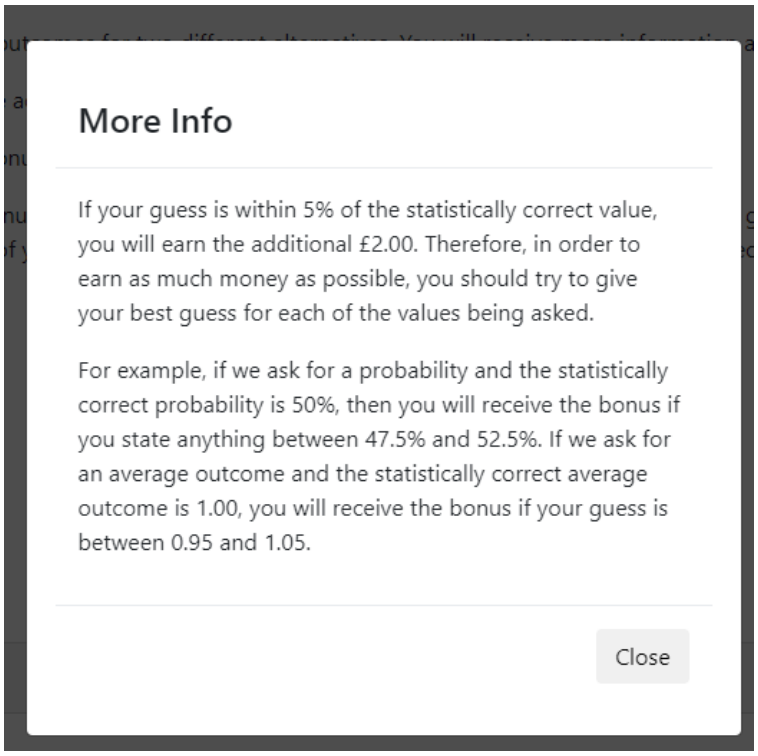


Figure IA17: Experiment A2: Belief Eliciation Screen. This figure shows the screen shown to participants in the beliefs group of Experiment A2 to ask for their beliefs about Alternatives 1 and 2.

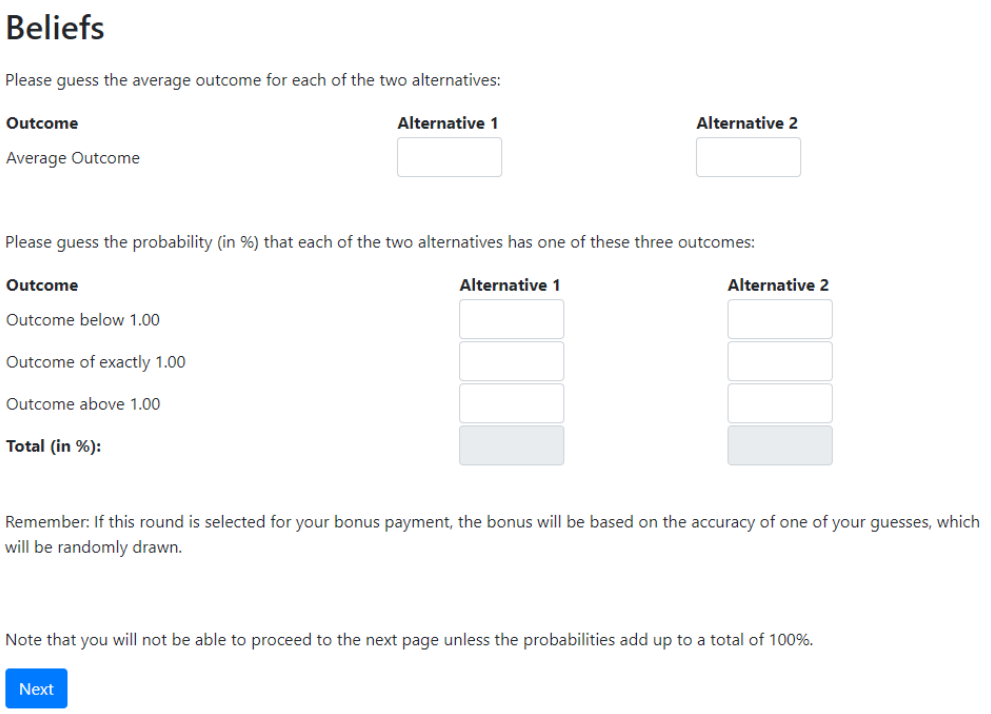


Figure IA18: Experiment A2: Information Screen. This figure shows the information screen shown to half of participants of Experiment A2 to facilitate learning about important characteristics of Alternatives 1 and 2. Both participants in the beliefs and the choice group are randomly selected to see this screen (or not) right before their main task (choice or beliefs).

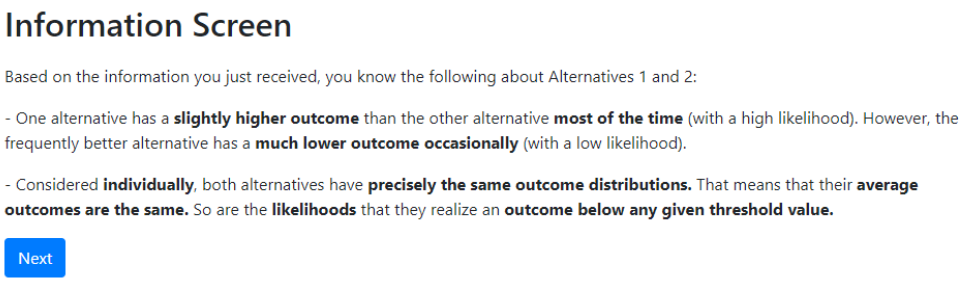


Figure IA19: Experiment A2: Recency or primacy effects. This figure shows the main outcome for participants in the choice sequential information group, split by when the extreme outcome pair was sampled. The split is done by quintile, from a bin with the first 20% to a bin with the last 20% of the sequence. The middle quintile includes one more observation (e.g., observations 5-7 in the short 11-state round, where other quintiles are 1-2, 3-4, 8-9, and 10-11). Point estimates are fractions of participants in the choice group that select the frequently-outperforming (but infrequently much worse) alternative. 95% confidence intervals are displayed for each point estimate, showing that all point estimates are significantly higher than 50%.

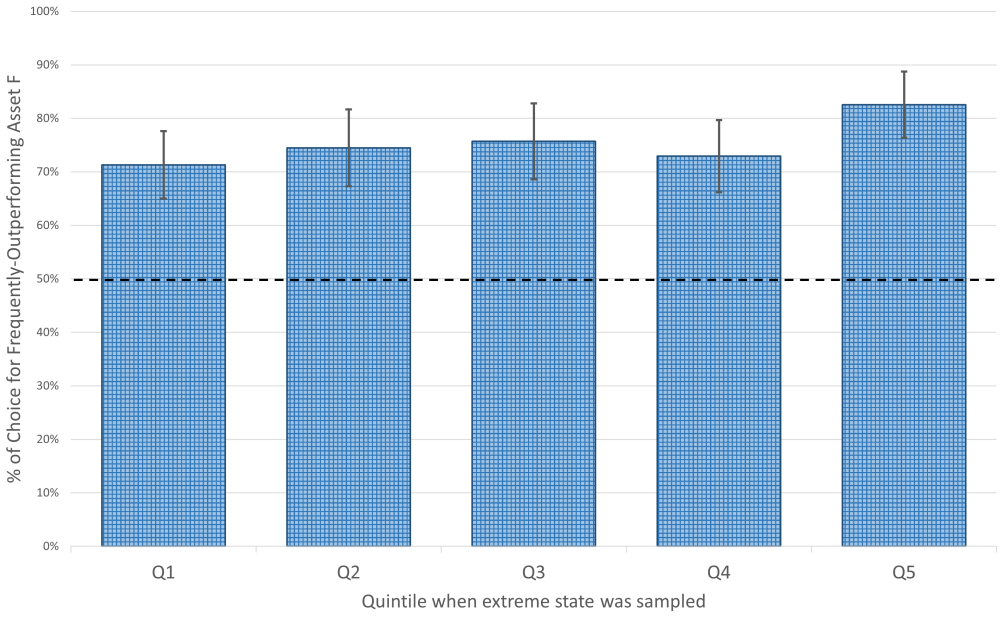
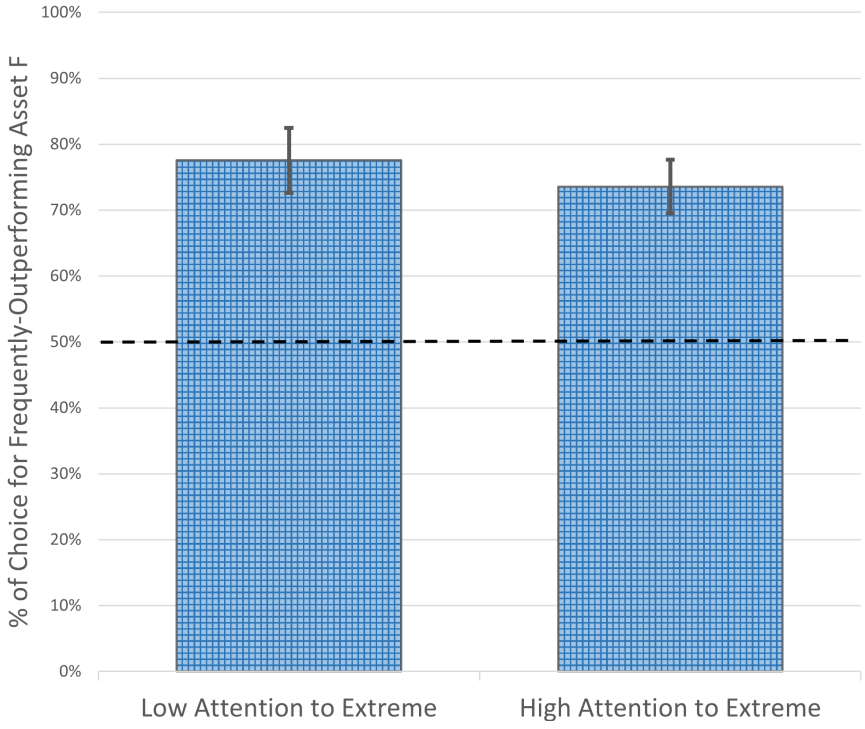


Figure IA20: Experiment A2: Split conditional on viewing-time of extreme outcome pair. This figure shows the main outcome for participants in the choice sequential information group, split by whether their viewing time for the extreme outcome pair was higher or lower than their average viewing time for the moderate outcome pairs. The result for the 63% (37%) of participants with a viewing time for the extreme pair above (below) their average viewing time for the moderate pairs is shown on the right (left). In both cases, a significantly higher fraction of participants selects the frequently-outperforming alternative.



C Additional Tables

Table IA1: Summary statistics. This table reports summary statistics for our Baseline Experiment and the theory-testing experiments, which systematically manipulate memory and attention. 'Age' is the participant's age in years and 'Female' is an indicator for the participant's sex. Participants self-assess their willingness to take risks on a scale from 1 (take no risks) to 5 (take large risks to achieve a significant gain). Self-assessed statistical knowledge is based on a scale from 1 (poor) to 5 (excellent). The student indicator is from Prolific. Participants state, on a scale from 1 (never) to 5 (very often), whether they consider what could happen if they made another decision (ex-ante regret) or could have happened if they make another decision (ex-post regret).

Variable	Mean	Std.Dev.	p10	p50	p90
Panel A: Baseline Experiment: 785 participants (Investors, 73% from the UK, others from the US)					
Age	38.68	12.91	24	36	58
Female	0.50	0.50	0	1	1
Willingness to take risks (1-5)	3.23	0.98	2	3	4
Self-assessed stat. knowledge (1-5)	2.71	1.02	1	3	4
Student	0.12	0.33	0	0	1
Considers regret ex-ante (1-5)	4.00	0.87	3	4	5
Considers regret ex-post (1-5)	3.70	1.02	2	4	5
Time for full experiment (seconds)	885.08	448.47	486	755	1418
Panel B: Theory-Testing Experiments: 600 participants (Investors, 74% from the UK, others from the US)					
Age	37.00	11.90	24	34	55
Female	0.50	0.50	0	0	1
Willingness to take risks (1-5)	3.31	1.00	2	3	5
Self-assessed stat. knowledge (1-5)	2.78	1.05	1	3	4
Student	0.20	0.40	0	0	1
Considers regret ex-ante (1-5)	4.04	0.87	3	4	5
Considers regret ex-post (1-5)	3.89	0.89	3	4	5
Time for full experiment (seconds)	781.76	401.07	388	689	1312

Table IA2: Interference Experiment – Joint Distribution. This table shows a description of the joint distribution, which participants in the "Interference" condition sampled. All 11 situations are equally likely and participants view a representative sample in random order, just like in the Baseline Experiment. The two assets that are available in the choice exhibit the same marginal outcome distributions, but are first-order stochastically dominated by a new third asset, the decoy company, which is sampled to interfer with memory formation. In the 'Interference (Unlikely-Outperforming)' variation, the decoy company is similar in color and industry to the 'Unlikely-Outperforming' asset (Tech Company 2). Interference in memory formation should thus improve the perceived attractiveness of the unlikely-outperforming asset, which is chosen by hardly anyone in the Baseline Experiment. In the 'Frequently-Outperforming' variation, the decoy company is similar in color and industry to the 'Frequently-Outperforming' asset (Oil Company 2). Inference in memory formation should thus confirm the perceived attractiveness of the frequently-outperforming asset, which is chosen by most participants in the Baseline Experiment. Note that the marginal distributions of the two available assets are slightly different from the Baseline Experiment. The key difference is that outcomes are on an irregular grid to have more digits, which is necessary to avoid systematic differences relative to the decoy in the "roundness" of outcomes. Marginal distributions are still exactly the same between the two assets on the choice menu (the last two columns) and the frequently-outperforming asset (last column) still outperforms moderately in 10 out of 11 situations, while underperforming by a large margin in 1 situation.

	Decoy Company	**Unlikely-Outperforming**	**Frequently-Outperforming**
Interference (Unlikely-Outperforming)	Tech Company 1 light blue	Tech Company 2 dark blue	Oil Company 1 dark orange
Interference (Frequently-Outperforming)	Oil Company 1 light orange	Tech Company 1 dark blue	Oil Company 2 dark orange
Situation 1	\$0.51	\$0.05	\$0.30
Situation 2	\$0.73	\$0.30	\$0.50
Situation 3	\$0.94	\$0.50	\$0.70
Situation 4	\$1.10	\$0.70	\$0.88
Situation 5	\$1.26	\$0.88	\$1.05
Situation 6	\$1.53	\$1.05	\$1.30
Situation 7	\$1.73	\$1.30	\$1.48
Situation 8	\$1.90	\$1.48	\$1.68
Situation 9	\$2.13	\$1.68	\$1.88
Situation 10	\$2.28	\$1.88	\$2.07
Situation 11	\$2.28	\$2.07	\$0.05

Table IA3: Baseline Experiment – Attention to Extremes. In this table, we report regressions of viewing times per outcome pair within the sequential information group of the Baseline Experiment on a dummy for the extreme outcome pair. Viewing time is measured in seconds (Specifications (1) to (3)) or log(seconds) (Specifications (4) to (6)). We split up all sampled outcome pairs into those sampled from short (11-pair) rounds versus long (51-pair) rounds in Specifications (2)-(3) and (5)-(6). *t*-statistics are based on participant-clustered standard errors. */**/*** indicates statistical significance at the 10%/5%/1% level.

	Seconds			Log(Seconds)		
	All	Short	Long	All	Short	Long
$I_{\text{Extreme Pair}}$	1.780***	1.554***	2.009***	0.552***	0.446***	0.654***
	6.58	3.81	6.17	18.85	14.13	14.69
Constant	1.867***	2.855***	1.665***	0.154***	0.556***	-0.007***
	10.91	77.07	260.73	6.52	193.36	-8.08
R^2	12.36%	21.43%	13.38%	60.15%	71.09%	61.18%
N	25,048	4,444	20,604	25,032	4,443	20,589
Participant fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Round fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observation number fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Long round fixed effects	Yes	No	No	Yes	No	No
Clustered by participant	Yes	Yes	Yes	Yes	Yes	Yes

Table IA4: Baseline Experiment – Choices and Viewing Times. In this table, we test whether longer viewing times for particular outcome pairs are associated with a choice that overweights these outcome pairs in the sequential information environment of the Baseline Experiment. We report regressions of a dummy indicating a choice for the frequently-outperforming asset minus 0.5 on a viewing time weighted average outcome difference between the unlikely-outperforming (U) and the frequently-outperforming (F) alternative, based on participant-round specific viewing times per outcome pair. The weighted outcome pair difference is standardized to have a zero mean and standard-deviation of 1 over all participant-rounds. *t*-statistics are based on participant-clustered standard errors. */**/*** indicates statistical significance at the 10%/5%/1% level.

Standardized viewing-time weighted $\overline{X^U - X^F}$	0.024
	(1.33)
Constant	0.349
	(19.59)
R^2	0.44%
N	404

Table IA5: Baseline Experiment – Choices. In this table, we report average choices in our Baseline Experiment. We report the average fraction of participants selecting each of the alternatives. See Section D.1 for a more detailed description of the different conditions. t-statistics are based on participant-clustered standard errors whenever both rounds are included. */**/*** indicates statistical significance at the 10%/5%/1% level.

Condition	Choice Asset U	Choice Asset F	Difference versus 50%	t-test	N	Strength of Preference (1-7)
All						
Sequential Information	16.46%	83.54%	33.54%***	23.28	808	5.84
Simultaneous Information	57.22%	42.78%	-7.22%***	-3.45	762	5.00
Choice Group						
Sequential Information	17.32%	82.68%	32.68%***	15.71	410	5.76
Simultaneous Information	54.64%	45.36%	-4.64%	-1.62	388	5.03
Beliefs Group						
Sequential Information	15.58%	84.42%	34.42%***	17.26	398	5.93
Simultaneous Information	59.89%	40.11%	-9.89%***	-3.24	374	4.97
Short round (11 outcome pairs)						
Sequential Information	15.10%	84.90%	34.90%***	19.57	404	5.72
Simultaneous Information	58.01%	41.99%	-8.01%***	-3.16	381	4.98
Long round (51 outcome pairs)						
Sequential Information	17.82%	82.18%	32.18%***	16.88	404	5.96
Simultaneous Information	56.43%	43.57%	-6.43%**	-2.53	381	5.02
Round 2 choice, when the round 1 choice was for the frequently-outperforming asset						
Sequential Information	11.71%	88.29%	38.29%***	21.70	333	5.87
Simultaneous Information	31.79%	68.21%	18.21%***	4.79	151	4.87
Round 2 choice, when the round 1 choice was for the unlikely-outperforming asset						
Sequential Information	32.39%	67.61%	17.61%***	3.15	71	5.69
Simultaneous Information	68.70%	31.30%	-18.70%***	-6.10	230	4.93

Table IA6: Baseline Experiment – Beliefs. In this table, we report average beliefs in our Baseline Experiment. Panel A reports the average outcome expected for each of the alternatives. Panel B reports the average probability of below/above-median outcomes for each of the alternatives. The t-statistics in Panel B are from a difference-in-difference estimate between the below- and above-median probability assessments. See Section D.1 for a more detailed description of the different conditions. t-statistics are based on participant-clustered standard errors whenever both rounds are included. */**/*** indicates statistical significance at the 10%/5%/1% level.

Condition	Beliefs Asset U	Beliefs Asset F	Difference (-in-Diff.)	t-test	N
Panel A: Expected outcome (0-2)					
All					
Sequential Information	1.02	1.16	0.15***	9.10	808
Simultaneous Information	1.08	1.01	-0.07***	-3.76	762
Choice Group					
Sequential Information	1.01	1.19	0.18***	7.33	410
Simultaneous Information	1.08	1.05	-0.03	-1.32	388
Beliefs Group					
Sequential Information	1.02	1.13	0.11***	5.47	398
Simultaneous Information	1.08	0.98	-0.10***	-3.69	374
Short round (11 outcome pairs)					
Sequential Information	1.00	1.17	0.16***	7.34	404
Simultaneous Information	1.07	1.00	-0.07***	-2.83	381
Long round (51 outcome pairs)					
Sequential Information	1.03	1.16	0.13***	6.52	404
Simultaneous Information	1.09	1.02	-0.07***	-2.67	381
Panel B: Risk assessment (below/above-median probabilities)					
All					
Sequential Information – P($x < 1$)	47.82%	37.48%	-20.44%***	-13.57	808
Sequential Information – P($x > 1$)	40.06%	50.16%			
Simultaneous Information – P($x < 1$)	39.53%	42.87%	6.07%***	4.04	762
Simultaneous Information – P($x > 1$)	44.64%	41.91%			
Choice Group					
Sequential Information – P($x < 1$)	48.43%	36.41%	-23.91%***	-10.24	410
Sequential Information – P($x > 1$)	38.88%	50.78%			
Simultaneous Information – P($x < 1$)	40.67%	42.50%	3.89%**	1.97	388
Simultaneous Information – P($x > 1$)	45.24%	43.19%			
Beliefs Group					
Sequential Information – P($x < 1$)	47.19%	38.59%	-16.85%***	-9.07	398
Sequential Information – P($x > 1$)	41.28%	49.53%			
Simultaneous Information – P($x < 1$)	38.35%	43.26%	8.34%***	3.67	374
Simultaneous Information – P($x > 1$)	44.02%	40.59%			
Short round (11 outcome pairs)					
Sequential Information – P($x < 1$)	48.40%	35.04%	-25.87%***	-12.58	404
Sequential Information – P($x > 1$)	38.24%	50.76%			
Simultaneous Information – P($x < 1$)	39.09%	42.63%	6.46%***	3.46	381
Simultaneous Information – P($x > 1$)	44.27%	41.35%			
Long round (51 outcome pairs)					
Sequential Information – P($x < 1$)	47.23%	39.92%	-15.00%***	-8.62	404
Sequential Information – P($x > 1$)	41.88%	49.57%			
Simultaneous Information – P($x < 1$)	39.96%	43.11%	5.68%***	3.12	381
Simultaneous Information – P($x > 1$)	45.01%	42.48%			

Table IA7: Baseline Experiment – Beliefs and Choices. In this table, we report average choices in our Baseline Experiment, split based on the beliefs about the expected outcome and risk of alternatives. We report the average fraction of participants selecting each of the alternatives. See Section D.1 for a more detailed description of the different conditions. t-statistics are based on participant-clustered standard errors. */**/*** indicates statistical significance at the 10%/5%/1% level.

Condition	Choice Asset U	Choice Asset F	Difference versus 50%	t-test	N
Believe frequently-outperforming asset has higher expected outcome					
Sequential Information	7.63%	92.37%	42.37%***	35.58	498
Simultaneous Information	42.04%	57.96%	7.96%**	2.01	157
Believe frequently-outperforming asset has lower expected outcome					
Sequential Information	29.83%	70.17%	20.17%***	5.91	181
Simultaneous Information	76.44%	23.56%	-26.44%***	-8.97	208
Believe frequently-outperforming asset has same expected outcome					
Sequential Information	31.78%	68.22%	18.22%***	4.43	129
Simultaneous Information	53.15%	46.85%	-3.15%	-1.26	397
Believe frequently-outperforming asset has lower risk					
Sequential Information	6.68%	93.32%	43.32%***	38.52	494
Simultaneous Information	46.15%	53.85%	3.85%	0.92	143
Believe frequently-outperforming asset has higher risk					
Sequential Information	43.61%	56.39%	6.39%	1.48	133
Simultaneous Information	75.20%	24.80%	-25.20%***	-9.28	254
Believe frequently-outperforming asset has same risk					
Sequential Information	23.20%	76.80%	26.80%***	8.52	181
Simultaneous Information	49.04%	50.96%	0.96%	0.37	365

Table IA8: Baseline Experiment – Recency and Primacy Effects in Choices. In this table, we report the average fraction of choices for the frequently-outperforming asset in our Baseline Experiment's sequential information environment, conditional on whether extreme outcome pairs where sampled in each of the five quintiles from the first to the last 20% of the sequence. The middle quintile includes one more observation (e.g., observations 5-7 in the short 11-state round, where other quintiles are 1-2, 3-4, 8-9, and 10-11). We report the average fraction of participants selecting each of the alternatives. t-statistics are based on participant-clustered standard errors whenever both rounds are included. */**/*** indicates statistical significance at the 10%/5%/1% level.

	Q1	Q2	Q3	Q4	Q5
Lower bound (95%-CI)	81.34%	78.85%	79.35%	74.51%	73.11%
Point estimate	86.06%***	84.30%***	85.21%***	80.99%***	79.72%***
Upper bound (95%-CI)	90.78%	89.75%	91.07%	87.46%	86.33%
Standard-error	2.41%	2.78%	2.99%	3.30%	3.37%
t-statistic	14.98	12.33	11.78	9.38	8.81
N	208	172	142	142	143

Table IA9: Summary statistics. This table reports summary statistics for Experiments A1 and A2. 'Age' is the participant's age in years and 'Female' is an indicator for the participant's sex. Participants self-assess their willingness to take risks on a scale from 1 (take no risks) to 5 (take large risks to achieve a significant gain). Self-assessed statistical knowledge is based on a scale from 1 (poor) to 5 (excellent). The student indicator is from Prolific. Participants state, on a scale from 1 (never) to 5 (very often), whether they consider what could happen if they made another decision (ex-ante regret) or could have happened if they make another decision (ex-post regret).

Variable	Mean	Std.Dev.	p10	p50	p90
Experiment A1: 574 participants (81% from the UK, others from the US)					
Age	32.68	10.70	21	30	49
Female	0.41	0.49	0	0	1
Willingness to take risks (1-5)	3.06	0.98	2	3	4
Self-assessed stat. knowledge (1-5)	2.63	1.05	1	3	4
Student	0.19	0.40	0	0	1
Considers regret ex-ante (1-5)	4.04	0.93	3	4	5
Considers regret ex-post (1-5)	3.91	1.02	2	4	5
Time for full experiment (seconds)	434.37	270.63	230	360	690
Experiment A2: 1520 participants (78% from the UK, others from the US)					
Age	36.09	11.86	22	34	53
Female	0.43	0.49	0	0	1
Willingness to take risks (1-5)	3.12	1.06	2	3	4
Self-assessed stat. knowledge (1-5)	2.69	1.08	1	3	4
Student	0.16	0.36	0	0	1
Considers regret ex-ante (1-5)	4.14	0.83	3	4	5
Considers regret ex-post (1-5)	3.89	0.98	2	4	5
Time for full experiment (seconds)	500.14	339.84	234	401	839

Table IA10: Experiment A1 – Choices. In this table, we report average choices in Experiment A1. We report the average fraction of participants selecting each of the assets. See Section D.2 for a more detailed description of the different conditions. t-statistics are based on participant-clustered standard errors whenever both rounds are included. */**/*** indicates statistical significance at the 10%/5%/1% level.

Condition	Choice Asset U	Choice Asset F	Difference versus 50%	t-test	N
All					
Sequential Information	27.05%	72.95%	22.95%***	9.81	414
Simultaneous Information Table	49.19%	50.81%	0.81%	0.29	370
Simultaneous Information Verbal	49.73%	50.27%	0.27%	0.09	364
In Levels					
Sequential Information	27.14%	72.86%	22.86%***	7.04	210
Simultaneous Information Table	52.72%	47.28%	-2.72%	-0.68	184
Simultaneous Information Verbal	50.56%	49.44%	-0.56%	-0.13	178
In Differences					
Sequential Information	26.96%	73.04%	23.04%***	6.80	204
Simultaneous Information Table	45.70%	54.30%	4.30%	1.09	186
Simultaneous Information Verbal	48.92%	51.08%	1.08%	0.27	186
Short round (11 outcome pairs)					
Sequential Information	23.19%	76.81%	26.81%***	9.12	207
Simultaneous Information Table	44.32%	55.68%	5.68%	1.55	185
Simultaneous Information Verbal	51.65%	48.35%	-1.65%	-0.44	182
Long round (101 outcome pairs)					
Sequential Information	30.92%	69.08%	19.08%***	5.93	207
Simultaneous Information Table	54.05%	45.95%	-4.05%	-1.10	185
Simultaneous Information Verbal	47.80%	52.20%	2.20%	0.59	182
Round 2 choice, when the round 1 choice was for frequently-outperforming					
Sequential Information	23.18%	76.82%	26.82%***	7.78	151
Simultaneous Information Table	37.50%	62.50%	12.50%**	2.41	88
Simultaneous Information Verbal	40.43%	59.57%	9.57%*	1.88	94
Round 2 choice, when the round 1 choice was for unlikely-outperforming					
Sequential Information	37.50%	62.50%	12.50%*	1.91	56
Simultaneous Information Table	53.61%	46.39%	-3.61%	-0.71	97
Simultaneous Information Verbal	62.50%	37.50%	-12.50%**	-2.41	88

Table IA11: Experiment A1 – Beliefs. In this table, we report average beliefs in Experiment A1. Panel A reports the average outcome expected for each of the assets. Panel A reports the average risk score given for each of the assets (from 'risk-free' at 1 to 'very risky' at 7). See Section D.2 for a more detailed description of the different conditions. t-statistics are based on participant-clustered standard errors whenever both rounds are included. */**/*** indicates statistical significance at the 10%/5%/1% level.

Condition	Beliefs Asset U	Beliefs Asset F	Difference	t-test	N
Panel A: Expected outcome (0-2)					
All					
Sequential Information	0.99	1.14	0.15***	6.04	414
Simultaneous Information Table	1.10	1.05	-0.05	-1.56	370
Simultaneous Information Verbal	0.97	0.95	-0.02	-0.64	364
In Levels					
Sequential Information	0.99	1.15	0.16***	4.87	210
Simultaneous Information Table	1.05	1.02	-0.03	-1.03	184
Simultaneous Information Verbal	0.95	0.95	0.00	0.05	178
In Differences					
Sequential Information	0.99	1.12	0.14***	3.72	204
Simultaneous Information Table	1.14	1.08	-0.06	-1.19	186
Simultaneous Information Verbal	0.98	0.94	-0.05	-0.98	186
Short round (11 outcome pairs)					
Sequential Information	0.94	1.15	0.21***	5.56	207
Simultaneous Information Table	1.10	1.06	-0.04	-1.26	185
Simultaneous Information Verbal	0.96	0.96	0.00	0.05	182
Long round (101 outcome pairs)					
Sequential Information	1.04	1.12	0.09***	3.23	207
Simultaneous Information Table	1.09	1.04	-0.05	-1.20	185
Simultaneous Information Verbal	0.97	0.93	-0.05	-1.05	182
Panel B: Risk assessment (1-7)					
All					
Sequential Information	4.17	3.48	-0.69***	-5.04	414
Simultaneous Information Table	3.81	3.86	0.05	0.48	370
Simultaneous Information Verbal	3.98	4.05	0.07	0.49	364
In Levels					
Sequential Information	4.15	3.36	-0.79***	-4.01	210
Simultaneous Information Table	3.89	3.91	0.02	0.12	184
Simultaneous Information Verbal	3.91	4.06	0.15	0.70	178
In Differences					
Sequential Information	4.20	3.61	-0.59***	-3.09	204
Simultaneous Information Table	3.72	3.81	0.09	0.50	186
Simultaneous Information Verbal	4.04	4.04	0.00	0.00	186
Short round (11 outcome pairs)					
Sequential Information	4.56	3.57	-0.99***	-5.41	207
Simultaneous Information Table	3.94	3.77	-0.17	-1.15	185
Simultaneous Information Verbal	3.92	4.18	0.25	1.27	182
Long round (101 outcome pairs)					
Sequential Information	3.79	3.39	-0.40**	-2.45	207
Simultaneous Information Table	3.68	3.95	0.27*	1.90	185
Simultaneous Information Verbal	4.03	3.93	-0.10	-0.53	182

Table IA12: Experiment A1 – Beliefs and Choices. In this table, we report average choices in Experiment A1, split based on the beliefs about the expected outcome and risk of alternatives. We report the average fraction of participants selecting each of the assets. See Section D.2 for a more detailed description of the different conditions. t-statistics are based on participant-clustered standard errors. */**/*** indicates statistical significance at the 10%/5%/1% level.

Condition	Choice Asset U	Choice Asset F	Difference versus 50%	t-test	N
Believe that frequent outperformer has higher expected outcome					
Sequential Information	12.83%	87.17%	37.17%***	17.46	304
Simultaneous Information Table	16.36%	83.64%	33.64%***	9.29	110
Simultaneous Information Verbal	24.31%	75.69%	25.69%***	5.83	144
Believe that frequent outperformer has lower expected outcome					
Sequential Information	80.00%	20.00%	-30.00%***	-6.08	70
Simultaneous Information Table	84.76%	15.24%	-34.76%***	-9.69	105
Simultaneous Information Verbal	81.25%	18.75%	-31.25%***	-8.23	112
Believe that frequent outperformer has same expected outcome					
Sequential Information	42.50%	57.50%	7.50%	0.51	40
Simultaneous Information Table	48.39%	51.61%	1.61%	0.38	155
Simultaneous Information Verbal	50.93%	49.07%	-0.93%	-0.27	108
Believe that frequent outperformer has lower risk					
Sequential Information	13.40%	86.60%	36.60%***	13.28	194
Simultaneous Information Table	26.60%	73.40%	23.40%***	4.99	94
Simultaneous Information Verbal	29.10%	70.90%	20.90%***	4.80	134
Believe that frequent outperformer has higher risk					
Sequential Information	46.28%	53.72%	3.72%	0.61	121
Simultaneous Information Table	71.68%	28.32%	-21.68%***	-4.39	113
Simultaneous Information Verbal	69.13%	30.87%	-19.13%***	-4.64	149
Believe that frequent outperformer has same risk					
Sequential Information	30.30%	69.70%	19.70%***	4.06	99
Simultaneous Information Table	46.63%	53.37%	3.37%	0.82	163
Simultaneous Information Verbal	48.15%	51.85%	1.85%	0.22	81

Table IA13: Experiment A2 – Choices. In this table, we report average choices in Experiment A2. We report the average fraction of participants selecting each of the assets. See Section D.3 for a more detailed description of the different conditions. t-statistics are based on participant-clustered standard errors whenever both rounds are included. */**/*** indicates statistical significance at the 10%/5%/1% level.

Condition	Choice Asset U	Choice Asset F	Difference versus 50%	t-test	N	Strength of Preference (1-7)
All						
Sequential Information	24.94%	75.06%	25.06%***	14.98	810	5.42
Simultaneous Information	50.14%	49.86%	-0.14%	-0.08	694	4.44
Without Info Screen						
Sequential Information	23.62%	76.38%	26.38%***	12.05	398	5.58
Simultaneous Information	53.47%	46.53%	-3.47%	-1.27	346	4.54
With Info Screen						
Sequential Information	26.21%	73.79%	23.79%***	9.43	412	5.27
Simultaneous Information	46.84%	53.16%	3.16%	1.18	348	4.34
Short round (11 outcome pairs)						
Sequential Information	22.96%	77.04%	27.04%***	12.92	405	5.29
Simultaneous Information	52.74%	47.26%	-2.74%	-1.02	347	4.39
Long round (101 outcome pairs)						
Sequential Information	26.91%	73.09%	23.09%***	10.46	405	5.56
Simultaneous Information	47.55%	52.45%	2.45%	0.91	347	4.49
Round 2 choice, when the round 1 choice was for frequently-outperforming						
Sequential Information	18.12%	81.88%	31.88%***	14.26	289	5.56
Simultaneous Information	49.71%	50.29%	0.29%	0.08	175	4.54
Round 2 choice, when the round 1 choice was for unlikely-outperforming						
Sequential Information	38.32%	61.68%	11.68%**	2.47	107	5.64
Simultaneous Information	51.74%	48.26%	-1.74%	-0.46	172	4.55

Table IA14: Experiment A2 – Recency and Primacy Effects in Choices. In this table, we report the average fraction of choices for the frequently-outperforming asset in Experiment A2's sequential information environment, conditional on whether extreme outcome pairs where sampled in each of the five quintiles from the first to the last 20% of the sequence. The middle quintile includes one more observation (e.g., observations 5-7 in the short 11-state round, where other quintiles are 1-2, 3-4, 8-9, and 10-11). We report the average fraction of participants selecting each of the assets. See Section D.3 for a more detailed description of the different conditions. t-statistics are based on participant-clustered standard errors whenever both rounds are included. */**/*** indicates statistical significance at the 10%/5%/1% level.

	Q1	Q2	Q3	Q4	Q5
Lower bound (95%-CI)	64.99%	67.30%	68.60%	66.19%	76.39%
Point estimate	71.29%***	74.48%***	75.69%***	72.94%***	82.55%***
Upper bound (95%-CI)	77.58%	81.66%	82.78%	79.69%	88.72%
Standard-error	3.19%	3.63%	3.59%	3.42%	3.12%
t-statistic	6.67	6.74	7.16	6.71	10.43
N	202	145	144	170	149

Table IA15: Experiment A2 – Attention to Extremes. In this table, we report regressions of viewing times per outcome pair within the sequential information environment on a dummy for the extreme outcome pair. Viewing time is measured in seconds (Specifications (1) to (3)) or log-seconds (Specifications (4) to (6)). We focus on the short (11-pair) round in Specifications (2) and (5), and on the long (51-pair) round in Specifications (3) and (6). t-statistics are based on participant-clustered standard errors. */**/*** indicates statistical significance at the 10%/5%/1% level.

	Seconds			Log(Seconds)		
	All	Short	Long	All	Short	Long
$I_{\text{Extreme Pair}}$	1.146***	0.969***	1.325***	0.384***	0.329***	0.438***
	7.53	3.53	8.78	20.17	15.18	14.68
Constant	1.618***	2.967***	1.261***	0.022	0.655***	-0.192***
	16.81	118.99	844.38	1.45	332.99	-650.13
R^2	7.17%	25.41%	5.60%	63.06%	72.13%	60.70%
N	92,078	9,152	82,921	92,078	9,152	82,921
Participant fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Round fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observation number fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Long round fixed effects	Yes	No	No	Yes	No	No
Clustered by participant	Yes	Yes	Yes	Yes	Yes	Yes

Table IA16: Experiment A2 – Choices and Viewing Times. In this table, we test whether longer viewing times for particular outcome pairs are associated with a choice that overweights these outcome pairs in the sequential information environment. We report regressions of a dummy indicating a choice for the frequently-outperforming asset minus 0.5 on a viewing time weighted average outcome difference between the unlikelz-outperforming (U) and the frequently-outperforming (F) alternative, based on participant-round specific viewing times per outcome pair. The weighted outcome pair difference is standardized to have a zero mean and standard-deviation of 1 over all participant-rounds. As an illustration, for participants in Specification (1) where all rounds are aggregated, a one standard-deviation increase in the viewing time weighted outperformance of the unlikely-outperforming asset is associated with a 3.9 percentage point lower likelihood to select the frequently-outperforming asset, relative to a baseline likelihood of 75.3% (0.500+0.253). That is, a viewing time weighted performance difference 6.53 (=0.253/0.039) standard-deviation above the average would be required to equalize the likelihood of selecting either asset in the sequential information environment. Notably, the average weighted outcome difference is already positive, because extreme outcome pairs receive overproportional attention on average (see Table IA15), with more than 62% of participants taking more time to view the extreme outcome pair than the average moderate outcome pair. t-statistics are based on participant-clustered standard errors. */**/*** indicates statistical significance at the 10%/5%/1% level.

	(1) All	(2) Short	(3) Long
Standardized viewing-time weighted $\overline{X^U - X^F}$	-0.039**	-0.013	-0.061***
	(-2.47)	(-0.63)	(-2.89)
Constant	0.253***	0.271***	0.235***
	(15.20)	(12.93)	(10.72)
R^2	0.80%	0.10%	2.04%
N	810	405	405

Table IA17: Experiment A2 – Beliefs. In this table, we report average beliefs in Experiment A2. Panel A reports the average outcome expected for each of the assets. Panel B reports the average probability of below/above-median outcomes for each of the alternatives. The t-statistics in Panel B are from a difference-in-difference estimate between the below- and above-median probability assessments. See Section D.3 for a more detailed description of the different conditions. t-statistics are based on participant-clustered standard errors whenever both rounds are included. */**/*** indicates statistical significance at the 10%/5%/1% level.

Condition	Choice Asset U	Choice Asset F	Difference (-in-Diff.)	t-test	N
Panel A: Expected outcome (0-2)					
All					
Sequential Information	1.04	1.13	0.08***	5.46	796
Simultaneous Information	1.08	1.05	-0.03	-1.35	740
Without Info Screen					
Sequential Information	1.05	1.14	0.09***	4.31	400
Simultaneous Information	1.08	1.05	-0.03	-0.91	360
With Info Screen					
Sequential Information	1.04	1.11	0.07***	3.38	396
Simultaneous Information	1.08	1.05	-0.03	-1.00	380
Short round (11 outcome pairs)					
Sequential Information	1.03	1.17	0.14***	6.48	398
Simultaneous Information	1.08	1.06	-0.02	-0.78	370
Long round (101 outcome pairs)					
Sequential Information	1.06	1.08	0.02	1.05	398
Simultaneous Information	1.07	1.04	-0.04	-1.19	370
Panel B: Risk assessment (below/above-median probabilities)					
All					
Sequential Information – P($x < 1$)	44.73%	41.14%	-7.26%***	-6.36	796
Sequential Information – P($x > 1$)	42.84%	46.50%			
Simultaneous Information – P($x < 1$)	43.91%	43.94%	0.44%	0.40	740
Simultaneous Information – P($x > 1$)	40.92%	40.50%			
Without Info Screen					
Sequential Information – P($x < 1$)	44.07%	40.56%	-6.96%***	-4.57	400
Sequential Information – P($x > 1$)	43.57%	47.01%			
Simultaneous Information – P($x < 1$)	44.38%	44.77%	0.66%	0.37	360
Simultaneous Information – P($x > 1$)	40.13%	39.86%			
With Info Screen					
Sequential Information – P($x < 1$)	45.39%	41.72%	-7.56%***	-4.28	396
Sequential Information – P($x > 1$)	42.10%	45.99%			
Simultaneous Information – P($x < 1$)	43.47%	43.14%	0.24%	0.18	380
Simultaneous Information – P($x > 1$)	41.67%	41.11%			
Short round (11 outcome pairs)					
Sequential Information – P($x < 1$)	44.04%	37.93%	-11.69%***	-6.52	398
Sequential Information – P($x > 1$)	41.50%	47.09%			
Simultaneous Information – P($x < 1$)	42.37%	42.33%	-0.22%	-0.16	370
Simultaneous Information – P($x > 1$)	40.33%	40.51%			
Long round (101 outcome pairs)					
Sequential Information – P($x < 1$)	45.42%	44.34%	-2.83%**	-2.07	398
Sequential Information – P($x > 1$)	44.18%	45.92%			
Simultaneous Information – P($x < 1$)	45.46%	45.54%	1.11%	0.70	370
Simultaneous Information – P($x > 1$)	41.52%	40.49%			

D Discussion of Additional Results

We discuss additional figures and tables from Internet Appendices B and C in this appendix. Section D.1 contains a few other results for our Baseline Experiment. Sections D.2 and D.3 discuss the results from two additional experiments that explore variations in learning environments and the robustness of the effect of learning environments on beliefs.

D.1 Baseline Experiment – Other Results

We report main results for our Baseline Experiment in Tables 1 and 2 under (S1) and (D1), as well as Figure 2. This appendix discusses some additional results, in particular splits by the incentivization scheme (choice vs. beliefs),[23] results for the long 51-state round of the experiment and persistence in choices across rounds, as well as choices conditional on beliefs. The summary statistics for the Baseline Experiment are in Panel A of Table IA1.

D.1.1 Choices

We report other results on choices for our Baseline Experiment in Table IA5. The first subpanel "All" combines all participants and rounds in each of the two presentation formats, including not just the short 11-state distribution we focus on in the main paper, but also the long 51-state distribution participants saw in the other round. Again, participants significantly favor the frequently-outperforming asset after the sequential information environment (83.54%, $p < .01$), but the unlikely-outperforming asset after the simultaneous information environment (57.22%, $p < .01$).

In this table, we also report the average self-assessed strength of preference for the participant's choice. It is high (always above 4.8 on a scale from 1 to 7) in both learning environments but systematically higher after sequential information. Participants in the simultaneous information environment have a systematically weaker preference for their choice (averages around 5 out of 7, vs. 6 out of 7 for the sequential information environment), pointing at some being close to indifference.

When splitting out participants who first made their choice and were incentiviced for their choice (not their beliefs), results are virtually the same in the sequential information environment. The

[23] See Figures IA5, IA6, IA7, and IA8 for screenshots of the four variations of choice/beliefs elicitation screens in the sequential vs. simultaneous information environment.

result for the simultaneous information environment is a bit weaker (54.64%, $p > .10$), but not statistically significantly different and it points in the same direction. Accordingly, the simultaneous information result is stronger (59.89%, $p < .01$) in the group of participants who were first asked for their beliefs and incentivized for accurate beliefs (and not their choice).

When splitting out the long 51-state round, we find qualitatively the same results as for the short 11-state round reported in the main paper. Learning in the sequential information environment is followed by significantly more choices for the frequently-outperforming asset whereas people select the unlikely-outperforming asset significantly more often in the simultaneous information environment.

There is strong persistence between rounds in the simultaneous information environment. The majority of participants who select the unlikely-outperforming asset after the simultaneous information environment in round 1 (230 of 381) tend to also do so in round 2 (68.70%, $p < .01$). The minority participants who select the frequently-outperforming asset in the simultaneous information environment in round 1 (151 of 381) tend to also do so in round 2 (68.21%, $p < .01$). We do not find this predictable heterogeneity for participants in the sequential information environment. Even the small minority of participants who select the unlikely-outperforming asset in round 1 (71 of 404) tend to select the other asset in round 2 (67.61%, $p < .01$).

D.1.2 Beliefs

We report other results on beliefs for our Baseline Experiment in Table IA6. Overall, results are consistent with our Baseline Experiment finding in the main paper (short 11-state round). After learning in the sequential information environment, participants are more optimistic about the frequently-outperforming asset. After the simultaneous information environment, they are more optimistic about the unlikely-outperforming asset. That is, belief patterns are consistent with choices from Table IA5.

Incentivicing participants for beliefs does not lead to less biased beliefs. Relative optimism for the unlikely-outperforming asset in the simultaneous information environment does not decrease at all, but rather increases (driven by more accurate beliefs about the other, frequently-outperforming asset though).

Ex-ante, one could have expected the overestimation of the average to shrink significantly in

the long versus the short round. After all, frequent states are characterized by the frequently outperforming asset having a \$0.04 higher outcome instead of the \$0.20 higher outcome from the short round. Nevertheless, participants expect a \$0.13 outperformance by the frequently better asset on average (versus \$0.16 in the short round), which is more than the *maximum* outperformance of the frequently-outperforming asset across all 51 states.

D.1.3 Beliefs and Choices

Last we analyze choices conditional on beliefs. This is not quite as clean as other tests, because we are splitting by participants' arguably noisy estimates for their beliefs, which could introduce biases. Nevertheless, Table IA7 shows a strong correspondence between stated beliefs and choices. Participants who are more optimistic for an asset are much more likely to select that alternative. Of course, we ask each participant for both beliefs *and* choices, so that participants might be adjusting beliefs to be consistent with choices or vice versa. Experiment A2 rules out this channel by asking half of participants only for beliefs, and the other half only for a choice.

D.1.4 Recency or Primacy Effects

One could argue that participants under-weight outcomes in the sequential information environment because they under-sample the relevant states (Fox and Hadar, 2006). Infrequent extreme states might be particularly likely to not be drawn at all. Our experimental design rules out sampling error as a cause for choosing the frequently-outperforming asset by showing each participant a representative sample. There is no sampling error and therefore every sample includes exactly one extreme state in our experimental design. Additionally, the order of outcome pairs is randomized for each sample. Therefore, systematic order effects—like extreme states occurring systematically at the beginning or end of the sample—can also be ruled out as a driver of choices.

However, one could still argue that recency or primacy bias could amplify or attenuate results. We test this by comparing choices in our Baseline Experiment when the extreme state is sampled in early versus late sections of the sample sequence. Specifically, we split the sample sequence into five sections ("quintiles"), separate choices by when the extreme state was sampled, and re-estimate our main result for the five subgroups.[24] Table IA8 and Figure IA9 show that whatever quintile of the

[24] Both rounds in the Baseline Experiment have an odd number of outcome pairs, which makes calculating average

sample sequence includes the extreme state, participants heavily favor the frequently-outperforming asset, with more than 79% of participants choosing the asset in all five subgroups and confidence intervals never extending below 73%.[25] In summary, we can reject that recency or primacy bias play a major role for the effect of the sequential information environment on choices in our experiment.

D.2 Experiment A1 – Robustness with other Presentation Formats

In Experiment A1, we test variations of the simultaneous information environment (tabular simultaneous information without a graph, purely verbal simultaneous information without a graph) and a variation where participants view not just outcome levels for both assets but also the difference (which may be argued to make extreme outcome pairs more salient).

Experiment A1 is programmed in oTree (Chen et al., 2016). Like the experiments in the main paper, it is run online, using Prolific (Gupta et al., 2021) to recruit participants from the US/UK. However, we do not select the investor filter on Prolific, and the experiment does not include bot and attention check at the beginning, or a comprehension check before round 1. Summary statistics for participants in Experiments A1 are reported in Panel A of Table IA9. This experiment was run in November 2022 and preregistered on AsPredicted.org, see this link. Including a base fee, the average hourly salary for the experiment is more than $20.

The design of Experiment A1 is very similar to the Baseline Experiment, but there are some differences. We now list the essential differences. Detailed instructions are reported in a supplementary document provided under this link.

First, Experiment A1 always first asks participants to make their choice between the two risky alternatives. Participants are subsequently also asked to state their beliefs for the two alternatives, but beliefs are not incentivized. In short, we do not include the beliefs group from the Baseline Experiment.

Second, there are two simultaneous information environments In Experiment A1, which are both different from the Baseline Experiment's simultaneous information environment. The main difference is the absence of a graph so that there is hardly any room for visual salience to matter. Some

outperformance particularly easy in our distributions. We include the odd observation in the middle quintile. For the 11-state sample, the groups are then 1-2, 3-4, 5-7, 8-9, and 10-11.

[25] These analyses include not just the short 11-state but also the long 51-state samples, which allows for more power in such splits. Point estimates are qualitatively the same when the two distributions (long and short) are separately analyzed.

participants view a purely verbal simultaneous information format, similar to classical decision science research, see Figure IA10. Others view a tabular simultaneous information format, which conveniently lists all outcome pairs sorted by one of the two assets' outcomes, see Figure IA11.

Third, for both the sequential and simultaneous information environment, some participants see not just levels of outcome variables, but also differences. See Figures IA12 and IA13 for the verbal and tabular simultaneous information environments including differences, and Figure IA14 for the sequential information environment with differences. Figure IA15 shows the sequential information environment without differences, which is nearly identical to the Qualtrics version from our Baseline Experiment.

D.2.1 Choices

Table IA10 shows our main result for Experiment A1. As in our Baseline Experiment's results, we find that more than 2 out of 3 participants favor the frequently-outperforming asset after learning in the sequential information environment. While the effect is large, it is not quite as large as for the Baseline Experiment (72.95% instead of 83.54% select the frequently-outperforming asset), which may be due to the inclusion of language and bot checks, comprehension questions, as well as the investor filter in the Baseline Experiment. For participants in both simultaneous information environments, choices are statistically insignificantly different from 50% instead of being tilted towards the unlikely-outperforming alternative as in the Baseline Experiment. This may be due to the minor role for visual salience in Experiment A1's non-graphical formats.

Overall, we also find a highly significant choice gap in Experiment A1. In line with our model and bounded memory as the driving force after sequential information, participants favor the frequently-outperforming asset in the sequential information environment. The lack of visually salient extreme states in both the verbal and the tabular simultaneous information environments is consistent with choices that do not significantly differ from 50% in favor of either asset.

When splitting out participants who view not just levels of outcomes, but also differences, we observe qualitatively the same results. Similarly, there is no meaningful difference between the short (11-case) distribution in this experiment and the long (in this experiment 101-case) distribution. The very lengthy sequential sampling procedure for the 101 outcome pairs in the long round might explain a slightly weaker choice pattern in that subsample compared to the short distribution.

Like in our Baseline Experiment, we find choices in the simultaneous information environment to be persistently heterogeneous. While participants do not favor one or the other asset on average, it is possible to use round 1 to identify those who tend to select the frequently-outperforming or the unlikely-outperforming asset. This is not the case for the sequential information environment, where even the few participants who selected the unlikely-outperforming asset in round 1 are significantly more likely to select the frequently-outperforming asset in round 2.

D.2.2 Beliefs

As for the Baseline Experiment, we find that participants' beliefs are consistent with their choices. In the sequential information environment, they are more optimistic about the frequently-outperforming alternative, see Table IA11, both for the expected outcome (Panel A) and their risk assessment (Panel B).[26] Like for choices, beliefs are not significantly different from the neutral response in the simultaneous information environment.

Table IA12 shows a strong correspondence between stated beliefs and choices. Participants who are more optimistic for an alternative tend to select that alternative. Of course, beliefs are not incentivized and elicited after the choice, which could drive participants to adjust their beliefs so that they are consistent with their choice. Experiment A2 rules out that these design elements of Experiment A1 drive the evidence for a beliefs-based mechanism.

D.3 Experiment A2 – Isolating Biased Beliefs

In Experiment A2, we make sure that our results are not driven by spill-over effects of one task (e.g., the choice) to the incentivized main task (e.g., the beliefs elicitation). We either ask participants only to make an incentivized choice or only to state incentivized beliefs. Additionally, some participants are also presented with an information treatment (summarizing characteristics of the distribution of outcomes like equal marginal distributions) in an attempt to de-bias beliefs.

Experiment A2 is programmed in oTree (Chen et al., 2016). Like the experiments in the main paper, it is run online, using Prolific (Gupta et al., 2021) to recruit participants from the US/UK. In contrast to the Baseline Experiment, we do not select the investor filter on Prolific, and the

[26] In contrast to our Baseline Experiment's more involved probability estimates, we had participants score how risky each alternative is on a simple scale from 1 to 7.

experiment does not include bot and attention check at the beginning, or a comprehension check before round 1. Summary statistics for participants in Experiments A2 are reported in Panel B of Table IA9. This experiment was run in March 2023 and preregistered on AsPredicted.org, see this link. Including a base fee, the average hourly salary for the experiment is more than $20.

The design of Experiment A2 is very similar to the Baseline Experiment and Experiment A1. However, there are some differences. We now list the essential differences. Detailed instructions are reported in a supplementary document provided under this link.

First, to reduce the number of variations compared to Experiment A1, we drop the difference formats from Experiment A1, which showed the difference between outcomes on top of the levels of outcomes for each situation. We also drop the tabular simultaneous information environment. We still use Experiment A1's verbal simultaneous information environment, which is different from the Baseline Experiment (see Section D.2).

Second, in contrast to Experiment A1, half of participants are only asked to state their beliefs, without any mention of a choice between the two risky alternatives. Accurate beliefs are incentivized, see Figure IA16, like in our Baseline Experiment. Consistent with our Baseline, we also elicit beliefs about the probability distribution, on top of beliefs about the expected outcome, instead of just a risk score like in Experiment A1, see Figure IA17. The other half of participants are only asked to make a choice, without any mention of beliefs. They are incentivized for their choice, as in Experiment A1. This avoids any influence of this secondary task on the main task.

Third, half of participants in each learning environment (sequential and simultaneous information) and task (beliefs and choice) additionally see an information screen, which verbally describes key features of the joint distribution, such as equal marginal distributions and the existence of an infrequent but extreme outcome pair, see Figure IA18. One could argue that such a description should de-bias beliefs about marginal distributions, which could then be used to analyze the causal effect of beliefs on choices.

D.3.1 Choices

Table IA13 confirms our main result from Experiment A1. The sequential information environment leads to a strong attraction towards the frequently-outperforming alternative. There is a highly significant choice gap, with choices based on simultaneous information not deviating signif-

icantly from 50%. The information screen does not seem to affect choices significantly, similar to the short vs. long distribution variation. We replicate our results from Experiment A1, although beliefs are never mentioned to participants in the choice group.

Consistent with our Baseline Experiment and Experiment A1, there is persistent heterogeneity in choices in the simultaneous information environment, though this persistence is less significant compared to the other experiments. Participants who select the unlikely-outperforming asset in the simultaneous information environment in round 1 tend to also do so in round 2, and vice versa for participants who select the frequently-outperforming asset in round 1. Again, we do not find this predictable heterogeneity for participants in the sequential information environment.

In contrast to Experiment A1, we collect the order of each sample and viewing times in Experiment A2, so that we can also run analyses on recency and primacy bias, as well as attention effects as in our Baseline Experiment (see Sections 4.3 in the main paper and D.1.4 in the Appendix).[27]

Table IA14 and Figure IA19 report the recency/primacy bias result for Experiment A2. As for our Baseline Experiment, we find that participants strongly favor the frequently-outperforming asset in the sequential information environment, whether they sampled the extreme state early or late in the sequence. In contrast to the Baseline, we find a slightly stronger effect when the extreme state is sampled late. This kind of "reverse recency bias" might be caused by participants speeding up and overlooking the extreme state after sampling for a while, particularly in the more lengthy 101-state distribution, which requires drawing 101 outcome pairs.

In line with our Baseline Experiment, we again find that sampling the extreme outcome pair leads to a significant increase in viewing times. Table IA15 shows the viewing time regressions for Experiment A2 and indicates that viewing times for the average draw increase from 1.62 seconds for a moderate pair to 2.76 seconds for an extreme pair, by 1.15 seconds ($p < .01$). So, extreme outcome pairs are attention-grabbing, but like in our Baseline Experiment we do not find that increased average attention towards the extreme pair changes average choices.

Figure IA20 splits participants into the many who pay overproportional versus the few who pay underproportional attention towards the extreme outcome pair. The main result does not show a meaningful difference. In Table IA16, we report viewing time regressions, which do show

[27] These results can be compared to our Baseline Experiment but not to Experiment A1, as we did not collect the order of outcome pairs and viewing times in Experiment A1.

statistically significant effects. Like in our Baseline Experiment, these are economically too small to meaningfully affect the main result, even for participants who pay many standard-deviations more attention towards the extreme outcome pair than the average participant.

D.3.2 Beliefs

As in our Baseline Experiment and Experiment A1, beliefs remain overoptimistic for the frequently better alternative in the sequential information environment (see Table IA17). This is despite participants only stating beliefs (not making a choice) and being incentivized for accurate beliefs. Like for choices, beliefs are not significantly different from the neutral response in the simultaneous information environment.

We find that the information intervention is ineffective in debiasing beliefs. This is consistent with choices in Table IA13 not reacting to the information treatment. More generally, it is consistent with the small effect non-graphical (verbal or tabular) simultaneous information environments in Experiments A1 and A2 have on choices and beliefs.