Optimal Memory with Sequential Learning: Signals or Posterior Beliefs

Collin Raymond* Lars Wittrock†

March 19, 2024

Abstract

Agents with memory constraints must make trade-offs as to what information to remember over time. They may choose to remember signals, and form posterior beliefs only when asked; or alternatively, they may only remember posterior beliefs, and neglect to remember past signals. We demonstrate that memory constrained agents who can flexibly and optimally choose what information to remember will alter their choices in response to changes in the decision-making environment. When there is more uncertainty about which states are relevant for a decision, or when there are fewer signals, agents will tend to remember signals, and only form posteriors from signals when required. In contrast, when there is little uncertainty about the decision-relevant states or many signals, agents choose to remember posterior beliefs over the relevant states, and neglect to remember signals.

Keywords: Sequential learning, information processing, memory constraint

JEL Classification: D83, D91, C91

*SC Johnson College of Business, Cornell University, USA.

†Department of Micro and Public Economics, Maastricht University, The Netherlands.

The experiment in this paper was pre-registered and can be accessed here: https://aspredicted.org/NQD_R74.

1 Introduction

When individuals make decisions that depend on the dynamic acquisition of information, they typically sequentially observe partially informative signals about some underlying states of the world, knowing that at some point they must stop and make a decision. This framework has been applied throughout economics, e.g. choosing how much data to acquire in sequential testing (Wald, 1945), deciding when to stop when engaging in search with learning (Rothschild, 1974), or deciding which of many risky options to choose in a multi-armed bandit problem (Bergemann and Valimaki, 2006).

Although traditionally economics has tended to assume that decision-makers can perfectly remember past information at zero cost, it is well documented that individuals have limited and imperfect memory. For example, it is well known that individuals have bounds on their ability to remember sequences of digits (Miller, 1956), informational chunks (Cowan, 2010), and verbal units (Oberauer et al., 2016) — for a survey see Ericsson and Kintsch (1995). A the same time, it appears individuals are at least partially aware of their memory bounds and try to allocate memory optimally (e.g. Brady et al., 2009; Da Silveira et al., 2020), including improving memory when incentives increase (Wieth and Burns, 2006; Spaniol et al., 2014).

Such issues of bounded memory become extremely important under sequential information acquisition, because decision-makers must choose what, and how much, to remember over time. A standard approach to modeling decision-makers in these environments assumes that at the beginning of each period the decision-maker has a prior belief, observes a signal, and updates the prior to a posterior, which becomes the prior for the following period. The typical assumption is that the agent perfectly remembers their posterior belief at the end of every period. In contrast, a distinct intuition is that decision-makers begin with an initial prior, and at the end of every period know that initial prior, and all past signals they have observed. Whenever they are prompted to make a decision, they combine the initial prior with the past signals that they remember to form a posterior belief. However, across periods this posterior belief is not remembered — each time the agent needs to make a decision they combine the initial prior with the set of remembered signals.

Of course, if decision-makers have perfect, costless memory and are Bayesian, these two cognitive processes are equivalent: the initial prior updated with the set of signals observed up until time t gives the same final belief over states as a sequential updating procedure where each period the decision-maker updates last period's posterior with the current period's signal, to form this period's posterior.

However, if agents have limited memory, these processes may not generate equivalent final beliefs and observable behaviors. Within economics, there have been numerous approaches modelling cognitively limited agents who either keep track of posteriors after every signal (e.g. Wilson, 2014) or agents who keep track of signals, and only when asked to make a decision combine those signals into a posterior (e.g. Bordalo et al., 2023). These papers demonstrate that limited memory, combined with either of these two cognitive processes, can generate important behavioral biases. Moreover, they tend to show that these two distinct cognitive processes can explain different kinds of biases.

Despite the fact that both processes have appealing intuitions, researchers have tended to assume decision-makers only use one of the two. To our knowledge, there have been limited attempts to explore which, of these two cognitive processes individuals actually employ, and if both, when they employ one or the other.

This paper develops a framework which allows us to predict and test, whether, and when, each of these two memory processes are being used by decision-makers. We show that individuals respond to the implicit incentives in the economic environment and trade off the costs of remembering different kinds of information — signals versus posteriors over states. We find that across two different treatments lab subjects respond in a way predicted by theory. Specifically, we focus on two central factors of the environment: (1) uncertainty about the decision-relevant dimension and (2) the number of signals individuals observe before making a decision. We find that environments with a large number of states, and uncertainty about which of them will be relevant for the decision, and few signals, lead decision-makers to focus on remembering signals, and updating beliefs only when necessary. In contrast, in situations with a clear decision-relevant dimension and many signals, decision-makers instead update posterior beliefs over states every period, and tend not to remember past signals. Our data can be rationalized by a simple framework of subjects optimally trading off the costs and benefits of remembering different pieces of information over time.

Such findings are in line with intuitions about real-world behavior. Consider the following two examples. First, a financial analyst may be tasked to assess whether a specific stock is likely to increase or decrease in value in the future. In this case the analyst is only interested in one particular dimension, the rise/decline in stock value of the particular company, when observing new information. Moreover, the analyst may observe a large number of signals when sifting through many different data sets one after the other. It is likely that in this setting, the analyst continuously updates a posterior belief regarding the likelihood of the stock value going up or not. On the other hand, consider the example of a detective piecing together different clues regarding the likelihood that different suspects are guilty. With only a small number of clues available and each clue providing some information about many different suspects, it is likely that the detective keeps track of all individual signals instead of updating a posterior belief for each suspect. In the end, the detective may form a posterior belief about a specific suspect.

In Section 2 we discuss the related literature on the topic of memory and sequential information processing. We provide evidence from the literature that people's memory is constraint and as a result people attempt to optimize their memory usage. Moreover, we categorize existing papers on sequential information processing by their assumption about memory of the decision-maker. Two assumptions are commonly made by existing theoretical papers: people are either assumed to keep track of only posterior beliefs or only individual signals. Categorizing existing papers reveals that the decision environments differ significantly between the different groups of models. In addition, existing experimental papers provide empirical support for both assumptions.

In Section 3, we develop a simple model that captures the essence of our approach. We consider a sequential learning environment where remembering information is costly but individuals can freely choose what information they would like to keep track of. There is a single rational decision-maker who receives multiple pieces of information over time. The decision-maker is constrained by a cost of remembering information. At each point in time the decision-maker can choose what information they want to remember until the next period. They can choose to remember aggregated information in the form of posterior beliefs, observed signals, or both. After the final period, the decision-maker is faced with a decision. We study the problem of optimal memory for the decision-maker and analyze the role of different decision environments on optimal memory. If individuals incur costs for remembering more information they should respond to changes in the environment by changing what kind of aggregated information they choose to remember. Specifically, we vary the number of decision-relevant dimensions and the number of signals. We find that 1) there exist scenarios in which a decision-maker never remembers any signals but instead keeps track of at least some posterior beliefs and 2) there exist other scenarios in which a decision-maker never keeps track of any posterior beliefs but remembers at least some signals.

In Section 4, we develop an experimental paradigm that allows us to test the key predictions of our model. In the experiment, subjects observe multiple pieces of information in a sequential manner before making a decision in the end. Importantly, our experimental design allows us to cleanly vary different factors of the environment. We compare the behavior of participants in two treatments. Specifically, we compare the behavior of subjects in an environment with relatively many signals and a clear question to the behavior in a setting with relatively few signals and multiple potential questions. The key challenge is to elicit what subjects keep track of in an incentive-compatible way while not affecting subjects' behavior in the task itself. We solve this challenge by eliciting the preferences for a posterior/signal memory question from each participant. This allows learning what participants keep track of while receiving new information. Finally, subjects are faced with a cognitively demanding distraction task after each signal to increase the memory load.

In Section 5 we provide the results from the experiment. We find robust differences between the two treatment conditions, showing that the learning environment clearly influences what people remember. The direction of the difference provides strong evidence in support of the theoretical predictions. The majority of subjects keep track of posterior beliefs in an environment with many signals but only a single decision-relevant dimension while the majority of subjects remember individual signals in an environment with few signals and multiple decision-relevant dimension. In addition, we analyze the accuracy of reported beliefs and recalled signals across the two treatments. Those people that kept track of posterior beliefs in each treatment report similar beliefs on average, and those people that kept track of individual signals recalled a specific signal with similar accuracy. Lastly, we provide an overview of subject's behavior in the distraction task.

Finally, Section 6 concludes with a discussion of the findings.

Our contribution to the literature on memory, information acquisition, and Bayesian updating is three-fold. First, we view our work providing a more careful understanding of what decision-makers choose to remember, and how this responds to changes in the environment. We extend work that shows that individuals alter how much they want to remember in response to changes in the economic environment and we demonstrate that they change the qualitative kind of information they choose to remember. Second, our approach unifies several strands of the literature in behavioral and experimental economics, which have typically either assumed individuals can keep track of a summary "posterior" (or some other sufficient statistic) or they can keep track of past signals, but not both. We show that both approaches are valid, depending on the environment under consideration. Third, we provide a new experimental framework for more closely understanding the impact of economic environments on memory. Our framework allows us to vary the cognitive costs associated with different memory strategies, and so is distinct from existing approaches which typically vary the incentives (i.e. benefits) of remembering.

Of course our paper also abstracts away from many important details, both in our model and experimental design. For example, we focus on two extreme memory strategies — we construct environments where we believe that a decision-maker will likely focus on remembering only posteriors or only signals. Of course, in the real world, they are likely to attempt to remember some of both. We also suppose that there is no cost to updating a belief. In reality, combining signals with posteriors is a cognitively demanding task, and likely one that an agent would like to do as few times as possible.

2 Literature

Human memory is an extensively studied topic in psychology and, more recently, economics. Kahana (2012) provides a thorough review of the foundations of human memory. Our paper most directly relates to the concept of 'recall'. We, in particular, focus on the fact that humans tend to have limited recall, and that the accuracy of recall seems to be influenced by incentives. Perhaps the most widely known paper on limited recall is Miller (1956), who found that recall memory is limited and on average, people can recall 7 distinct chunks of information (i.e. numbers) from working memory. The precise limitations of working memory is not entirely settled, e.g., in a different setting, verbal chunks, Cowan (2010) finds that the capacity limit of working memory is reached with only 3 to 5 different items (Oberauer et al. (2016) provides a recent review of empirical work on limited memory, with views towards understanding the form of the constraints). Ericsson and Kintsch (1995) provides an overview of evidence that long term memory is similarly capacity constrained.

There is also a literature that suggests that humans alter their use of memory in the face of incentives. Many studies have suggested (see, e.g. Miller (1956) and Brady et al. (2009) for classic and recent takes on this idea), that individuals can re-code or compress information to make better use of limited memory, while some papers (such as Brady et al., 2009) show that this can occur as a response to heightened incentives. Building on these insights, authors such as Bays and Husain (2008); Ma et al. (2014) suggest that memory should be viewed as a limited resource that one can flexibly allocate between different tasks.

Within economics and related fields, a fairly diverse set of approaches have been taken in order to try and model, as well as test, limited memory. Most papers typically make one of two assumptions. The first is that a decision-maker keeps track of a posterior belief (or single summary statistic) regarding the state of the world, updates this in response to each new signal, and then has no further need to recall signals. The second is that a decision-maker attempts to keep track of all individual signals. Only when asked to make a decision do they attempt to aggregate those signals into a belief over states of the world, which they then use to make a decision. We turn to discusses papers using each approach in detail below. Table 1 in the appendix provides a summary. The first column of the Table lists the article, while the second categorizes the contribution as either theoretical (t) or empirical (e). The third column discusses the type of cognitive bounds placed on the the decision-maker, whether they be a finite set of memory states, the ability to only selectively or imperfectly recall memories, the decay of memory, or a cost of memory. The final column summarizes important details about the setting, including the number of states, signals, and what kind of decisions the decision-maker must take. The papers are categorized by their assumption regarding the memory of information over time, i.e. posterior beliefs or individual signals, or whether the decision is endogenous.

We first turn to papers that assume individuals keep track of posteriors (or a summary statistic) across periods. In early theoretical work Cover and Hellman (1970) propose a model of learning with finitely many memory states in the two-armed bandit problem. They assume the learner can only remember a summary statistic (e.g., a posterior) that can take on finite values. Wilson (2014) generalizes their model, and characterizes the optimal protocol. The memory constraint as formalized by Wilson is also applied in settings with dynamically changing states (Monte and Said, 2010; Chatterjee et al., 2022), arbitrary termination each period (Hu, 2023), endogenous learning (Kocer, 2010; Chatterjee and Hu, 2023), and strategic settings such as cheap talk games (Monte, 2005) and zero-sum games (Monte, 2013, 2014). In a similar vein, Dow (1991) considers an agent engaged in sequential search, but can only recall past summary information.

The second approach, that people remember individual signals and only form posteriors when needed to make decisions, originated in economics with Mullainathan (2002). In more recent work, papers such as Bordalo et al., Wachter and Kahana (2023), Enke et al. (2023), Graeber et al. (2022), Fudenberg et al. (2022), Neligh (2022) and Leung (2023) make a similar assumption that people recall past signals when asked to make a decision, albeit with very different assumptions about the form of recall. While many assume that that past signals are selectively recalled based on association or similarity, others operate under the assumption that memory decays, while some allow for both. A somewhat distinct literature on motivated beliefs, beginning with Bénabou and Tirole (2002); Bénabou and Tirole (2004), extended by Gottlieb (2014) and Chew et al. (2020), also assumes that individuals recall past signals rather than their previous posteriors.

We next turn to existing experimental work in economics and psychology that looks at individuals' ability to either recall past signals or provide beliefs about the state of the world in environments that feature dynamic information acquisition.

While there is a voluminous literature eliciting posterior beliefs over states of the world, a much smaller literature is explicitly devoted to looking at situations where signals are acquired sequentially. Although not explicitly motivated by concerns about limited memory, in a recent survey paper Benjamin (2019) pools data from eight existing studies where subjects were shown signals sequentially to examine how people update. Benjamin distinguishes two types of behavior, 'pooling' and 'acceptive'. Pooling behavior implies that whenever people observe a new signal they pool together all signals they have received up until that point and update their belief using their initial prior belief — implying that individuals should keep track of signals, not their posteriors, between periods. Acceptive behavior, on the other hand, implies that after each new signal people update their belief and this updated belief becomes the new prior belief — implying that individuals keep track of posteriors between periods. Benjamin finds no evidence for pooling behavior and therefore suggests that people are acceptive, i.e. they update their belief after every new signal. However, many of the studies used by Benjamin (2019) feature environments where our approach would predict that posteriors are relatively easy to keep track of — e.g., few decision-relevant states.

In contrast, other recent work, such as Enke et al. (2023), Graeber et al. (2022) and Bordalo et al. (2023), are motivated explicitly by models that assume individuals remember signals. They tend to find support for the conjecture that individuals carry around memories of past signals (i.e. evidence for what Benjamin (2019) would call pooling). Outside of economics, papers such as Shadlen and Shohamy (2016) also find that subjects seem to form beliefs about states of the world only when asked, and do so by sampling past signals from memory. That said, some of these settings, such as Bordalo et al. (2023), feature details that our model would predict should induce such behavior, such as subjects not knowing what they will be asked about in the future.

There are a few papers that are closer to our question: do we observe individuals trading off whether to remember signals or posteriors? Da Silveira et al. (2020) study the optimal memory of a decision-maker in a dynamic forecasting problem, and find that it is optimal to keep track of a single aggregate summary statistic in this setting. But their environment has features that make keeping track of signals relatively costly, and aggregate statistics relatively cheap. d'Acremont et al. (2013) find evidence that there are two distinct regions working in parallel when forming subjective beliefs. One region of the brain combines the signal with prior information to update a posterior belief while another region encodes the frequencies of observing signals.

3 Framework

3.1 Model

We next turn to providing a formal model that encapsulates the trade offs that a decision-maker with limited memory would face when trying to decide whether to remember posteriors over states or signals. We keep the model as simple as possible in order to accentuate the novel intuitions.

Environment As in most models involving decisions under uncertainty, we begin by describing the set of states. We assume a state space with a particular structure: there are D independent dimensions, and abusing notation slightly, indexed by $d = 1, ..., D$. Each dimension d can take on one of two realizations, 0 or 1. The state space is then denoted by $\Omega = \times_D d$ (with representative state ω which can be associated with a vector of length D with entries 0 or 1, corresponding to the realization of each of the dimensions). For each dimension d, there exists a bet $b_{d,k}$. A bet is a menu of two acts. The act $a_{d,0}$ pays off $L > 0$ if dimension d has realized value $k = 0$ and 0 otherwise; and the act $a_{d,1}$ pays off $L > 0$ if dimension d has realized value $k = 1$ and 0 otherwise. $\Delta(b)$ is a set of lotteries over bets, with generic lottery q. If a given bet is realized, the decision-maker must choose between the two acts in that bet. We denote $m(q)$ as the maximum weight applied to any dimension in distribution q; i.e. $m(q) = \max_{d \in D} q(d)$.

The decision-maker has a prior belief $\psi \in [0, 1]^D$ over realizations of each dimension, where ψ_d indicates the likelihood that dimension d has realization 1. This generates a prior over states Ω in the natural way. In each period $t = 1, ..., T$ the decision-maker observes an independently drawn signal s from the set of possible signals S. The Blackwell matrix $G_{\omega,s}$ gives the probability of observing signal s, conditional on state ω.

Decision problem. The decision-maker faces a decision problem q, L, T, where they will first observe T signals sequentially and then, at time period $T+1$, will face a bet b that is drawn from q.[1] L is the parameter that governs the size of the payoffs from the bet. Next the decision-maker chooses one of the two acts from the realized bet, then the state is realized. If the decision-maker chose the correct act from the bet they receive a payoff of L and otherwise 0.

Information processing and memory. Remembering information across periods comes at a cost to the decision-maker. We first describe how the agent can process information in any given period t.

At the beginning of each period t, the decision-maker has a set of state variables remembered from the previous period:

- the original prior belief $\psi \in [0, 1]^D$;[2]
- a set $R^t \subseteq D$ of dimensions where the belief from the prior period is inherited, and for each $d \in R$, an associated ρ^t_d, which indicated the (posterior) belief that $d = 1$.
- A set M^t of signals remembered from period $t-1$, which is a collection of pairs (m, τ), one for each τ between 1 and $t-1$, where $m \in S \cup \emptyset$

In each period t, after observing signal s^t, the decision-maker can first update their belief about different dimensions. We assume that updating beliefs is costless and the decision-maker therefore updates their beliefs about all dimensions.[3] We denote $\beta^t_d(\psi, R^t, M^t, s^t) \in [0, 1]$ as the Bayesian update of beliefs about dimension d at time t given the initial state variables, and the observed signal in this period.

The decision-maker chooses a strategy ζ to process information. A strategy ζ is a sequence of mappings, $\zeta^1, ..., \zeta^T$, where each ζ^t is a mapping from any possible ψ, R^t, M^t, s^t to a set of remembered beliefs, R^{t+1}, and signals, M^{t+1}, (see below for details). We assume that an agent always knows the strategy they are employing, and that the agent can commit to the ex-ante optimal strategy. As part of the strategy ζ, the decision-maker can then choose among several possible actions to remember the information for next period. We assume that remembering a chunk of information (either a particular signal or the belief about a particular dimension) incurs a fixed cost. We allow for memory costs to vary between beliefs about dimensions and signals. The cost for remembering a given dimension we denote as c_d, while the cost for remembering a given signal is c_s.

- The decision-maker can choose for any $d \subset D$, to remember $\beta^t_d(\psi, R^t, M^t, s^t)$. For each d that they choose to remember, they must pay a fixed cost c_d. Let $y^t_d(\phi, R^t, M^t, \zeta^t) \in \{0, 1\}$ be a binary variable

[1]Our results are qualitatively similar if if the agent, instead of knowing when they must take a decision, faced uncertainty about when the bet must be taken.

[2]Of course, it is not obvious that a decision-maker would be able to remember the original prior costlessly. One natural interpretation of our assumption is that ψ is simply whatever belief the agent would use as a prior, in the absence of remembering anything else. The substantive assumption is then that the agent always uses the same ψ regardless of the time period, or anything else they do remember. However, we view such an assumption as a natural starting point for our model.

[3]In Section 3.3 we relax this assumption and discuss the implications.

indicating whether the decision-maker remembers beliefs about dimension d at the end of time t. If dimension d is remembered then $d \in R^{t+1}$ and $\rho_d^{t+1} = \beta_d^t(\psi, R^t, M^t, s^t)$

- The decision-maker can choose which signals they wish to remember until the next period. Given M^t and s^t, the decision-maker can pay a cost c_s to remember any given element of $M^t \cup (s^t, t)$. If they choose to pay the cost for a given (m, τ) then $m_\tau^{t+1} = m_\tau^t$ with $1 \leq \tau \leq t-1$ or $m_\tau^{t+1} = s^t$ with $\tau = t$. If they don't pay the cost for a given τ then $m_\tau^{t+1} = \emptyset$. This implies that once a signal was forgotten it cannot be recovered in the future. Let $z_\tau^t(\phi, R^t, M^t, \zeta^t) \in \{0, 1\}$ be a binary variable indicating whether the decision-maker remembers signal τ at the end of time t. Taken together, this implies:

$$m_\tau^{t+1} = \begin{cases} \emptyset & \text{if } z_\tau^t = 0 \\ \begin{cases} m_\tau^t & \text{for } \tau \leq t-1 \\ s^t & \text{for } \tau = t \end{cases} & \text{if } z_\tau^t = 1 \end{cases}$$

Preferences and the Optimization Problem. We assume that the decision-maker has complete knowledge about the environment and the decision problem. We assume the decision-maker is risk neutral and when choosing the act the decision-maker acts as a Bayesian, with full knowledge of their optimal decision process, and chooses to maximize their expected payoff. To simplify matters, we also assume that there is no discounting.

In order to simplify our analysis we suppose the decision-maker has commitment power, and so commits to a ζ ex-ante. Given the prior ψ, and a strategy ζ the total cost of this strategy is given by:

$$C(\zeta) = \sum_{t=1}^{T} \sum_{d=1}^{D} c_d \cdot y_d^t(\zeta^t) + \sum_{t=1}^{T} \sum_{\tau=1}^{t} c_s \cdot z_\tau^t(\zeta^t) \tag{1}$$

In period $T+1$, given a realized bet $b_{d,k}$ and a belief $\beta_d^t(\phi, R^{T+1}, M^{T+1})$ the agent chooses between $a_{d,0}$ and $a_{d,1}$. The expected payoff is then

$$E[B(\zeta)] = E[\max\{a_{d,0}, a_{d,1}\} | \beta_d^t(\psi, R^{T+1}, M^{T+1})] \tag{2}$$

Thus, the optimization problem the decision-maker faces is then given by:

$$\max_{\zeta} E[B(\zeta)] - E[C(\zeta)]$$

3.2 Results

The model we have presented can give nuanced and complicated optimal memory rules depending on the circumstances. The goal in this paper is to highlight two relatively simple implications of the model, which don't require us to specify the entire optimal policy for all possible parameter combinations. We provide sufficient conditions on the environment so that either one of two things happen: in one set of environments the decision-maker will only focus on keeping track of posterior beliefs over states, and not remember signals between periods, while in other set of environments they will do the opposite. We do not claim (and in fact have numerous counterexamples) that these are the only two kinds of policies that a decision-maker would use. Rather we want to highlight the fact that the decision-maker responds to incentives in the environment, and then look for traces of these responses in the data.

We begin with the proposition that is somewhat easier to state. It says that, fixing a number of signals, we can make the bets important enough, and have enough dimensions and enough uncertainty about which dimension will be bet on, that that the decision-maker never updates their posteriors, but will remember some signals.

Proposition 1 *There exists a $\bar{T}$, $\bar{D}$, a $\bar{m(q)}$ and a $\bar{L}$ such that for all problems q, L, T, where $T \leq \bar{T}$, $D \geq \bar{D}$, $m(q) \leq \bar{m(q)}$ and $L \geq \bar{L}$, for every $t \leq T$*

- *for all t $R^t = \emptyset$, and*
- *for some $\tau \leq t-1$ and s^t $m^t_\tau = m^{t-1}_\tau$ and $m^t_t = s^t$.*

Proof:

In order to prove the result, set $\bar{T} = 1$. First we will show that if the agent, if they do not have the option of remembering posterior beliefs, but only signals, would want to remember the signal. Then we will show that the agent will want to remember the signal and not posterior beliefs.

1. The cost of remembering the signal is c_s. This means that the cost of remembering the single signal from $t = 1$ to $t = 2$ (the time when the bet is realized and chosen) is then c_s. Suppose the realized bet is on dimension $\hat{d}$. Consider a given history of signals h. Conditional on remembering the signal, and given the realized bet, denote the posterior belief used for betting as $B_{\hat{d}}(h)$. Given L, this pays off a value $Z_{\hat{d}}(h) = L \max\{B_{\hat{d}}(h), 1 - B_{\hat{d}}(h)\}$. Taking the expectation over the different signals and bets gives an expected benefit of $\sum_h \sum_d \sum Z_d q(d) f(h)$, where f is the ex-ante probability of history h. Notice that this is bounded below by $\frac{1}{2}L$, and is independent of the number of dimensions or the distribution over bets. We can set L large enough so that this exceeds c_s. Thus the agent would choose to remember the signals if that was the only possibility. The total utility from remembering signals is bounded below by $\frac{1}{2}L - c_s$.

2. Let the number of dimensions go to infinity. Let $m(q)$ go to 0. Then the benefit of remembering the posterior attached to any given state is bounded above by $m(q)L$. The total utility from remembering the posterior for this dimension (and not remembering ths signal) would be $m(q)L - c_d$. We can construct $m(q)$ small enough so that this is always negative. Thus, the agent would never want to remember the posteriors attached to any state. $\square$

Thus, if it is important to have as much information about the states of the world as possible, and there are relatively few signals, but a lot of uncertainty about what kind of bet will be given, then signals will be remembered, but states won't. Importantly, at least some signals will be remembered, but potentially not all.[4]

The next proposition is more nuanced. It says that that we can find a scenario where the decision-maker is certain enough about what bet they will face, but where there are enough signals, and the bets are important enough so that by the end of the sequential information acquisition process: (i) the decision-maker doesn't choose to remember any signals, but (ii) updates and remembers posteriors over a subset of the states of the world.

Proposition 2 *There exists $\bar{m(q)}$, $\bar{T}$ and $\bar{L}$ so that for all problems q, L, T where $m(q) \geq \bar{m(q)}$, $T \geq \bar{T}$, $L \geq \bar{L}$, there exists a $\tilde{T} < T$ such that for every $t \geq \tilde{T}$:*

- *$\tau \leq t$, $(m, \tau) = (\emptyset, \tau)$,*
- *for at least one d, $\rho_d^{t+1} = \beta_d^t(\psi, R^t, s^t)$ and $d \in R^{t+1}$ (unless $\rho_d^{t+1} = \psi_d$).*

Proof: We prove the result for $m(q) = 1$ so that the decision-maker knows for sure which bet they will face. Call the dimension that the bet will be on $\hat{d}$. This immediately implies that the decision-maker only cares about their posteriors over dimension $\hat{d}$. The proof happens in several steps. First we show that the decision-maker prefers to remember the relevant (the one attached to $\hat{d}$) posterior in all periods rather than some subset (including no periods), if signals are never remembered. Second we show that the decision-maker would find it worse to remember all signals than the relevant posterior in all periods. Third, we show that the agent, if they remember posteriors, need not remember the signals. Fourth, we show it is sub-optimal for the agent, once they begin to remember the posterior, to stop remembering the posterior in any future period. Last, we show that at some point it is cheaper to remember the posterior than the signals.

1. We will first compare a feasible strategy, remembering all the information from all periods (via posteriors), to a counterfactual strategy, which isn't actually feasible. We compute a strategy where the agent remembers the information acquired from a strict subset of periods $\mathbb{T}$, where the cost of remembering the information for a given period is $c = \min\{c_d, c_s\}$. This gives a total cost of remember of $c\mathbb{T}$. Notice that this cost isn't necessarily achievable, but it presents a useful counterfactual, as it is weakly lower than either the costs of remembering the information from $\mathbb{T}$ via either posteriors or signals.

   Denote the posterior for dimension $\hat{d}$ (since this is the only relevant dimension) after observing all signals given a history h as $B_{\hat{d}}(h)$. Denote the posterior after remembering a subset of signals from periods $\mathbb{T} \subset T$ as $B_{\hat{d}}^{\mathbb{T}}(h)$. Given L, the former pays off a value $Z_{\hat{d}}(h) = L\max\{B_{\hat{d}}(h), 1 - B_{\hat{d}}(h)\}$ and the latter $Z_{\hat{d}}^{\mathbb{T}} = L\max\{B_{\hat{d}}^{\mathbb{T}}(h), 1 - B_{\hat{d}}^{\mathbb{T}}(h)\}$. Taking expectations over histories (remember there is only a single dimension we are concerned about) we get an expected value of $\sum_h Z_{\hat{d}} f(h)$ (recall $f(h)$ denotes the distribution over histories from an ex-ante perspective) and $\sum_h Z_{\hat{d}}^{\mathbb{T}} f(h)$. Denote the

[4]This should not be surprising. Consider a sophisticated decision-maker who always faces a binary signal realization, either red or blue. By only remembering when red happens, they know that all "empty" memories are blue, and so can recall the entire sequence of memories. Moreover, there are values of L where "strong" signals would be remembered but not weak ones (which would all pool together as unremembered).

difference between these two, for a given $\mathbb{T}$ as ᛗ$(\mathbb{T}$ (recall ᛗ is runic m). Notice that ᛗ$(\mathbb{T}$ is positive and increasing (without bound) in L. This implies that for large enough L, ᛗ$(\mathbb{T} - c_d T - c\mathbb{T})$ is positive, implying that the decision-maker prefers to remember the posterior attached to $\hat{d}$ in all periods rather than the information from some subset of periods.

2. Next we show that the agent wouldn't want to remember all signals in all periods. If the agent remembers all signals, they would never remember any posteriors. The payoff is the same from remembering the posterior every periods, but the costs are bounded below by $c_s \frac{T(T+1)}{2}$, which for any fixed c_s, c_d must eventually be larger than $c_d T$ for a large enough T.

So we have now established that the agent prefers to remember information for all periods (via remembering posteriors) to remembering information from only a subset of periods.

Now we just need to show that there exists some t such that after t the agent will only remember posteriors and not remember signals.

3. If the agent remembers their posterior in Period t, they have no need to remember any signals that Period. Notice that all information relevant for the bet is contained in the posterior.

4. If the agent has remembered only the posterior in the past period, then they must remember the posterior this period, or they will lose information (unless the posterior is exactly equal to the initial prior $\psi^{\hat{d}}$ in which case the decision-maker acts "as if" they remembered the posterior, even though they didn't, which as we showed in the first step, is suboptimal.

5. Consider an agent in time period t, who currently remembers all signals up to time period t. The cost of remembering all signals to the next period is $c_s t$. The cost of remembering the posterior is c_d. For large enough t' the former is larger than the latter. Moreover, for all $t'' \geq t'$ it is also less costly to simply remember the posterior (which contains all information) than remember the signals (which also contain all information). Therefore, at time t, the agent will strictly prefer to remember the posterior for all periods moving forward,

Therefore, for t large enough, the agent will begin to remember the posterior, and only remember the posterior after that. Thus, we just need to ensure that $\hat{T} > t'$. $\square$

This proposition points out that if it is important to have as much information about the states of the world as possible, and there are relatively few states that matter compared to the total number of signals, then the decision-maker remembers posterior beliefs and not individual signals.

It is relatively clear that if the agent's posterior in any given period is equal to the initial prior, it is optimal to not remember the posterior, and simply start over the with prior. However, there are L's where it is optimal for the agent to remember the posterior so long as it is far enough from the initial prior, and otherwise forget the posterior, and pool on the initial prior.[5]

[5] This is because, if the current posterior is close enough to the initial prior, then the fix cost of remembering the posterior for a period outweighs the gain in information in moving from the prior to the posterior. Notice however, though that as the

Of course, the proposition does not say that the agent would only remember posteriors the entire process. It only says that by the end of the process the agent is only remembering posteriors. This is because if the cost of remembering a single signal is relatively small compared to the cost of remembering a posterior, the agent may want to initially only remember signals. Once there are enough signals such that the costs of remembering all the signals becomes prohibitive, the agent will aggregate all the signals into a posterior, and then remember posteriors forever after. Thus, it is key that there are enough signals, relative to the cost differential between remembering signals and posteriors, in order for the proposition to hold true.

One might wonder under what circumstances a stronger result would be true — where the decision-maker always remembers the posteriors across all periods. One situation is where there is a single dimension that will be bet on, and where the cost of remembering a posterior is the same as remembering a signal. More generally, albeit not surprisingly, we can find a small enough cost of cost of remembering a posterior relative to a signal (fixing all other parameters), so that an individual will only remember posteriors.

3.3 Extensions

One natural extension is to ask is whether, e.g., Proposition 1 could be extended to say that for any given T, we can find high enough incentives, and enough uncertainty about which question will be asked, so that only signals will be remembered. The answer, not surprisingly, is yes. Regardless of the uncertainty about the questions, so long as the benefit from guessing correctly is high enough, the agent wants to have all information available. However, the return to knowing about any given dimension is low enough when there is much uncertainty about what the bet is, so that the cost for remembering is lower for signals than posteriors. The answer, however, for the equivalent generalization for Proposition 2 (fixing a q) is different. This is because some signals might provide information about lots of potential bets. If these signals are sufficiently rare, it is better to simply remember them when they occur, rather than incorporating them into posteriors.

A second extension is to consider what happens if there is a fixed cost of Bayesian updating – every time the agent wants to combine a signal with a belief, they have to pay a cost. Again, this will not influence Proposition 2, but will change Proposition 1. In particular, this will also lead to "batch" processing of signals, where it could be that the decision-maker remembers both posteriors and signals for a time, and then turns all the signals into a posterior, forgets them, and then starts the process over.

A third extension would relax the assumption that either the decision-maker remembers for sure, or forgets for sure. In reality it is likely that a decision-maker can only choose a probability of remembering, where higher probabilities incur higher costs. In this case, the agent may have an incentive to not just remember current posteriors, but also past signals, as mis-remembered posteriors may be corrected by remembered signals. Moreover, an agent may even want to remember not just their most current posterior, but also past

cost of remembering a posterior goes to 0, the set of posteriors that are not remembered converges to the empty set. Of course, this may lead to the agent, as a sophisticated Bayesian updater, may actually learn something if they enter a period only remembering the initial prior (because they know their posterior beliefs from the previous period must be withing the set that wouldn't be remembered. This implies their Bayes update may not be the same as if their if their true beliefs are the initial prior.

posteriors.

4 Experimental Design

4.1 Design

Our experimental design is motivated by several objectives: 1) the setting should present clearly defined "chunks" of information to subjects in a sequential manner; 2) the setting should allow for easy variation of the environment (number of signals and decision-relevant dimensions); 3) the environment should be neutral to avoid motivated reasoning/memory in any form; and 4) the setting should allow for the elicitation of beliefs and the memory of signals. This leads to the following guessing task.

Guessing task. The guessing task has three dimensions, with two equally likely states in each dimension. Subjects receive noisy hints that are simultaneously informative for all three dimensions. The task is framed to subjects as the selection of three people that jointly determine the winning numbers for a local raffle. There are three teams, Team 'Parity', Team 'Round' and Team 'Size', with two members each. One member from each team is randomly selected with equal chance. Together the three selected people decide on multiple numbers in the range from 1 to 100. People from each of the three teams have different preferences over certain numbers. Members of Team 'Parity' like even/odd numbers, members of Team 'Round' like round/irregular numbers (i.e. divisible by 5 or not) and members of Team 'Size' like high/low numbers (i.e. greater than 50 or not). All three people makes sure that on average 75% of winning numbers are according to their preference. As an example, a hint may look as follows: 'The first winning number chosen by the three randomly selected people is: 12'. The complete description, as shown to subjects, can be found in Appendix B. Subjects are shown multiple signals, each followed by a short distraction task. The distraction task is explained in more detail below.

Treatments. We compare the behavior of subjects in two treatments with different environments. The two treatments differ in two ways: first, the number of signals that are shown to subjects and second, the number of potential questions subjects could be asked at the end. In the treatment *'posterior'* subjects are shown 14 signals and they are told that their task is to guess which person from Team 'Parity' was selected based on the hints they have seen. Subjects in the treatment *'signals'* are shown 4 signals. They are specifically told that their task is to answer one of three potential questions at the end. Any one of the three dimensions (i.e. teams) is equally likely to be selected as the question at the end. Participants of the experiment are randomly allocated to one of the two treatment conditions.

Elicitation. Eliciting whether subjects keep track of aggregated information in the form of posterior beliefs or individual signals is the key challenge with the experimental design. Different methods could be employed that focus on comparing ex-ante preference for difference questions or the ex-post performance of subjects. The elicitation procedure should fulfill several requirements: 1) elicitation should be incentivized; 2) the procedure should not influence the behavior of subjects during the guessing task and while seeing hints; 3) the elicitation method should not be influenced by the environment, i.e. if the same number of subjects keep track of posteriors/signals in two different treatments, the outcome should be the same; 4) subject's own (sub-conscious) choice of what to remember should directly affect the outcome of the elicitation method, i.e. the payment method, description of instructions, the level of the posterior or the signals themselves all should not influence the elicitation.

These objectives lead to the implementation of the following elicitation procedure. After observing roughly 75% of all signals, subjects are presented with a surprise question.[6] They have to choose which type of question they would prefer to answer to earn some additional monetary payoff. The two types of questions are either 'a question regarding the chance which person from Team 'Parity' was selected in the beginning' or 'a question regarding one of the numbers I have previously seen'. This question is purposefully not announced in the beginning such that the behavior of subjects during the main task is not influenced. Also, subjects are asked to make a decision before seeing all hints to avoid that subjects have already formed a posterior anticipating that they need to do so in order to answer the final question. With both types of questions subjects can earn \$8.00.

The benefit of this type of elicitation procedure is that it is easy to to explain to students and simple to implement. Subjects' behavior in the experiment should (almost) directly influence their question choice, i.e. subjects that remembered posterior throughout the experiment should choose the posterior question and subjects that remembered individual signals should choose the signal question. It could happen that subjects remembered individual signals and prefer to answer a question regarding the posterior.[7] However, switching in the other direction is not possible.

After subjects select the type of question they prefer to answer, they are immediately presented with this question. The implementation of the posterior question is comparable to many other belief elicitation tasks in the literature. Subjects have to select a percentage number from a slider regarding the relative likelihood of the two outcomes. Subjects are incentivized to report their belief accurately through a binarized scoring rule with rewards of \$8.00 or \$0.00. Subjects that instead chose for the signal question are asked about one randomly chosen number that was previously shown to them. In the language of Kahana (2012) this method could be considered 'cued recall' where the order of numbers acts as a cue for the recall of the signal. Subjects that correctly recall the signal earn \$8.00 while an incorrect response pays no additional bonus. Finally, after providing an answer to either the posterior or the signal memory question, subjects are asked to state their confidence (on a scale from 0 to 100) in the accuracy of their answer.

Memory load. After subjects observe a signal (i.e. at the end of every round), they are confronted with a distraction task, inspired by Deck et al. (2021). This task is purposefully designed to induce a memory load, making it harder for subjects to perfectly remember every detail of the experiment. The distraction task consists of two parts, a number memory question and multiple math questions. After explaining the task, an 8-digit number is shown to subjects for a few seconds. Then, they have 21 seconds to solve up to four math questions before being prompted to recall the 8-digit number. If subjects correctly recall the 8-digit

[6]Specifically, in the 'posterior' treatment this took place after 11 signals and in the 'signal' treatment after 3 signals.

[7]Although our theoretical model assumes away any cost of forming a posterior belief from a prior and signal, we recognize that in reality there is likely a cognitive cost associated with this action. This makes it less likely that subjects switch from remembering signals to answering a posterior question. In addition, this cost should be similar between the two different treatments, meaning that treatment differences are unaffected.

number they receive \$0.50. If subjects correctly answer one randomly selected math question they receive another \$0.50. Figure 4.1 provides an overview of the timeline of the experiment.

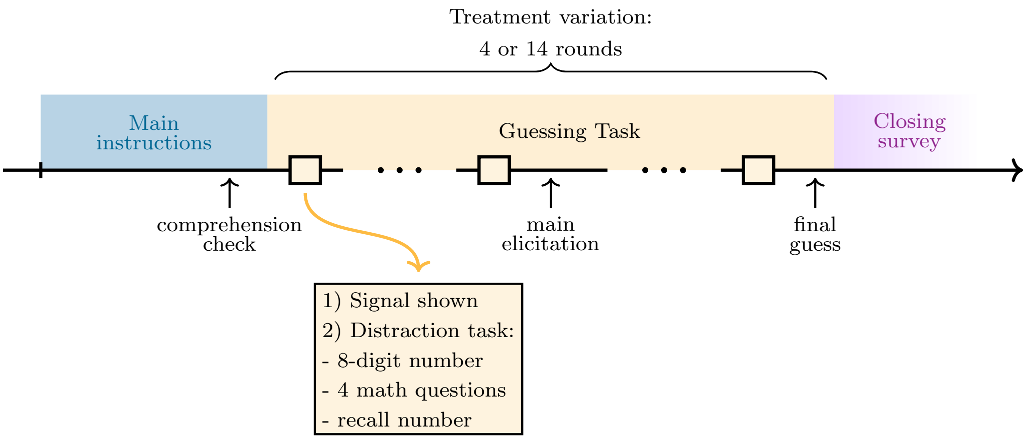

Figure 1: Timeline of the experiment

4.2 Hypothesis

The results from section 3.2 allow us to define a clear hypothesis regarding the behavior of subjects in the two treatments of the experiment. Two factors of the environments are varied in the two treatments, the number of signals and the number of decision-relevant dimensions. Proposition 2 states that, ceteris paribus, an environment with more signals before an eventual decision implies that more subjects should remember posterior beliefs rather than individual signals. Proposition 1 states that, ceteris paribus, an environment with fewer decision-relevant dimensions implies that more subjects should remember posterior beliefs rather than individual signals. In the two experimental treatments we simultaneously vary both factors.[8] The 'posterior' treatment has more signals and fewer decision-relevant dimensions than the 'signals' treatment condition. Therefore, the main hypothesis can be summarized as follows.

Hypothesis 1 In the 'posterior' treatment more subjects remember posterior beliefs (rather than individual signals) than in the 'signals' treatment.

This hypothesis and the analysis thereof presented in the following section were pre-registered.[9]

5 Results

This section presents the results from the experiment described above. The experiment was conducted at Cornell University in fall 2023. 100 subjects completed the experiment (49 in treatment 'posterior' and 51 in treatment 'signals'). Participants are students at Cornell University, mainly American, on average 25 years old, 2/3 are female and most study engineering, business/management, or other (social) sciences. On average, they took 39 minutes to complete the experiment and they earned \$12.59.

[8]This is done to create a larger potential treatment effect.

[9]The pre-registration can be accessed here: `https://aspredicted.org/NQD_R74`.

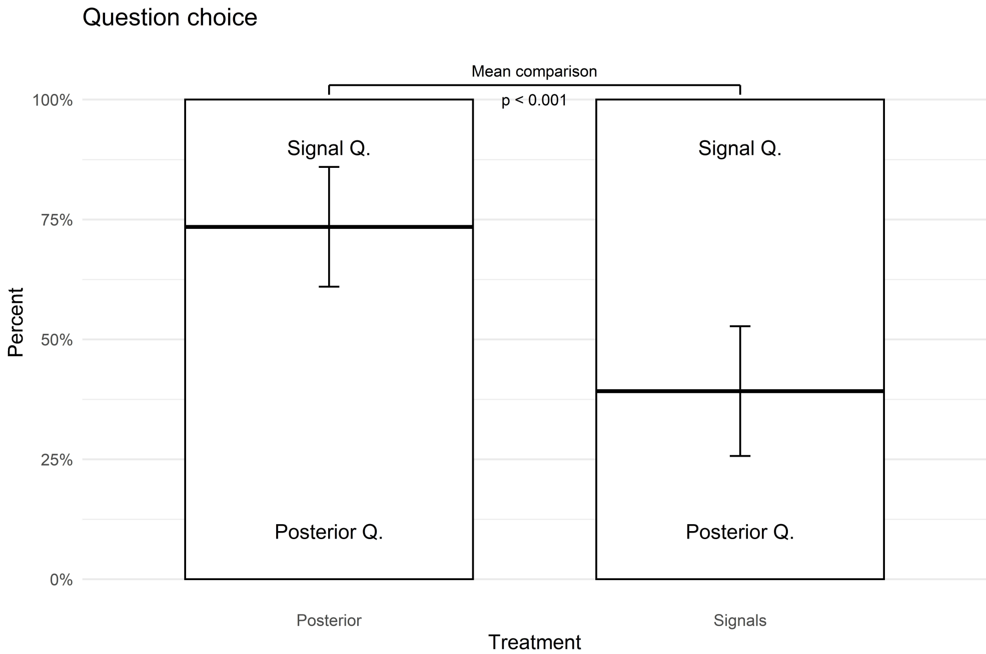

Figure 2: Question choice across the two treatments. Error bars show the 95% confidence interval around the depicted sample mean.

5.1 Main Hypothesis

The main question we investigate in this paper is whether people keep track of an aggregate posterior belief or individual signals when seeing multiple pieces of information over time. To answer this question we elicit subjects' preferences over two different types of questions. We compare two treatments that differ in the number of signals and the number of potential questions subjects may need to answer.

Figure 2 provides an overview of the frequency with which subjects choose the posterior/signal question for the two treatments. 73% of subjects choose the posterior question in the 'posterior' condition while only 39% of subjects choose the posterior question in the 'signal' condition. This difference is highly significant (p-value from chi-squared test: < 0.001), thus confirming Hypothesis 1. In addition, we estimate the following logistic regression:

$$\text{QuestionChoicePost}_i = \alpha + \beta_1 \cdot \text{TreatSignal}_i + X_i + \epsilon,$$

where QuestionChoicePost is a dummy variable that is equal to 1 for subjects that chose the posterior question, TreatSignal is a dummy variable indicating the treatment, and X_i is a set of additional explanatory variables. We find that, as hypothesized, β_1 is negative and significantly different from zero. The regression output is reported in Table 2 and the corresponding marginal effects of the different explanatory variables are reported in Table 3.

Result 1 In an environment with more signals and fewer decision-relevant dimensions, more subjects remember posterior beliefs (rather than individual signals).

We find strong evidence that the setting influences how people process sequential information. In an environment with many signals and a clearly defined question, people are significantly more likely to keep track of a posterior belief rather than individual signals. This implies that subjects rationally react to the environment when (unconsciously) choosing how to process information.

5.2 Additional Results

In the second part of the analysis we further analyze the accuracy of subject's answers to the posterior and signal memory question. We compare the accuracy and confidence of people who indicated that they remembered posterior beliefs, or individual signals respectively, across the two treatments.

5.2.1 Posterior

This section analyzes the behavior of all subjects who indicated that they remembered posterior beliefs. The sample is unbalanced between the two treatments as 36 subjects chose the posterior question in the 'posterior' treatment while only 20 subjects chose the posterior question in the 'signal' treatment. In both treatments, subjects are asked to indicate their posterior belief as a percentage between 0 and 100 by adjusting a slider (without default value).

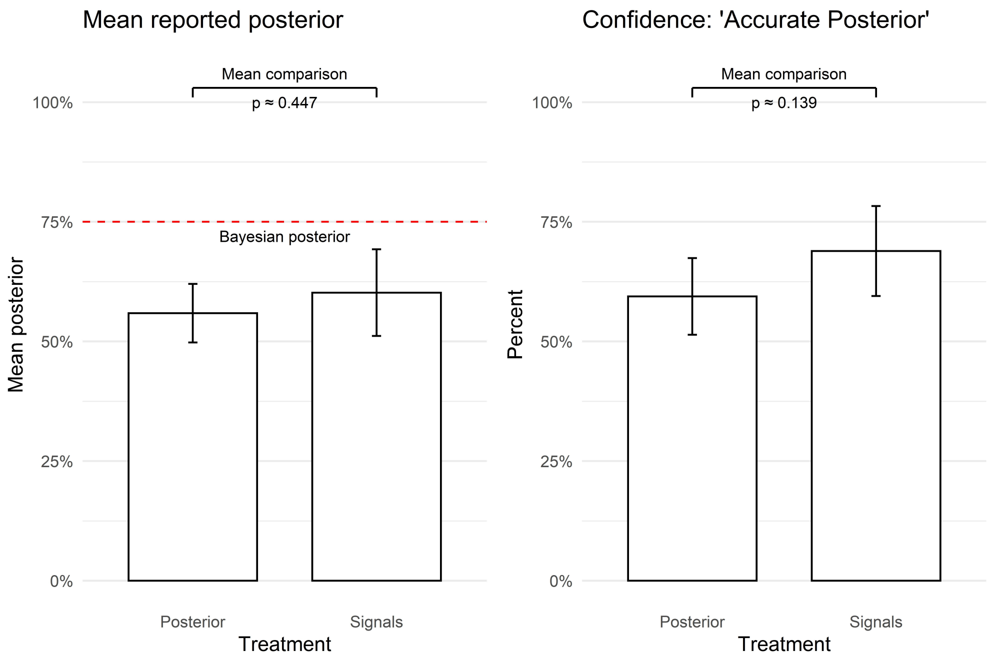

Figure 3: Average reported posterior beliefs and confidence by subjects who chose this question. The red dashed line indicates the Bayesian posterior given the 11 or 3 previous signals subjects have seen in the respective treatment.

Figure 3 shows the average reported posterior beliefs by subjects as well as the average level of confidence in the accuracy of the reported belief across the two treatments. The Bayesian belief report is identical across the two treatments at 75%. On average, reported beliefs do not significantly differ in the two treatments. In both treatments, the reported posterior beliefs are on average significantly more conservative than the

Bayesian posterior. Figure 5 in the appendix provides a more detailed overview of individual belief reports. Reported confidence regarding the accuracy of their belief report is also not significantly different on average in the two treatments.

In summary, for subjects that chose to keep track of posterior beliefs we do not observe differences in the accuracy of their (reported) beliefs or their confidence. It is important to note that these reports are only from the group of subjects that actively chose the posterior question. It is possible that the group of subjects that kept track of individual signals would hold different beliefs on average across the two treatments.

5.2.2 Signal Memory

This section analyzes the behavior of all subjects who indicated that they remembered individual signals. The sample is unbalanced between the two treatments as only 13 subjects chose the signal memory question in the 'posterior' treatment while 31 subjects chose the signal memory question in the 'signal' treatment. In both treatments, subjects are asked about one particular signal they have previously seen. Subjects are not provided with possible answers to choose from but need to freely recall the number between 1 and 100 that was shown to them before.

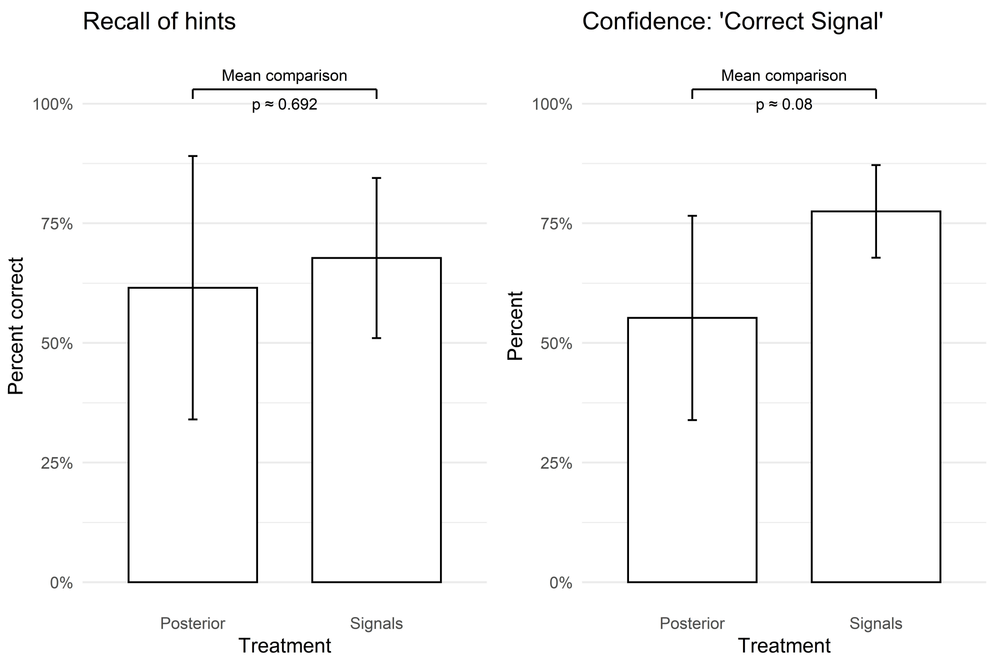

Figure 4: Percentage correct recall of the signal and confidence by subjects who chose this question.

Figure 4 shows the percentage of subjects that correctly recalled the signal, as well as the corresponding level of confidence across the two treatments. On average subjects recalled the signal with similar accuracy in the two treatments. This may be somewhat surprising as subjects in the 'posterior' treatment have seen in total 11 signals prior to answering the question, while subjects in the 'signals' treatment have only seen 3 signals. In both treatments, the majority of subjects correctly recalled the signal in question providing clear evidence that some subjects focus on carefully remembering all signals they observe. Reported confidence in the accuracy of their recall is slightly higher in the 'signals' treatment, though only significant at the 10% level. It is possible that a larger sample size would lead to a more robust difference.

In summary, in both treatments, the majority of subjects who chose the signal memory question recalls the specific signal correctly. While reported confidence is (potentially) higher in the 'signals' treatment, the actual recall accuracy is similar.

5.3 Distraction Task

After each signal, subjects are shown a distraction task consisting of a number memory question and multiple math question. This section provides an analysis of subjects' behavior in the two distraction tasks. The results show that subjects were fully engaged with the experiment and fatigue or other effects of time are unlikely to have impacted the treatment differences observed earlier.

We compare two aspects of the distractions task over time and across treatments: attempted solves and correct answers. First, we focus on attempted solves by subjects. A distraction task is considered to have been attempted if a subject provides any (potentially false) answer. Attempted solves are relevant as they provide an indication that the distraction task indeed created a mental load on subjects after each signal. Figure 6 in the appendix provides an overview of subjects' behavior. We find that nearly all subjects provide an answer for either the memory or the math tasks, and most for both. Moreover, subjects remain highly engaged with the distraction tasks until the end of the experiment and there are not differences in behavior across the two treatments. This provides evidence that fatigue and other undesired round effects are unlikely to have an effect on behavior in our main study.

Second, we analyze the actual performance of subjects in the distraction tasks. Figure 7 provides an overview of the number of correctly solved memory tasks per subject and Figure 8 provides an overview of the number of correctly solved math tasks per subject. We find that subjects differ substantially in their performance in the two distraction tasks. While few subjects answered no question in either the memory or the math task correctly, some subjects answered nearly all questions correctly. Again, it appears that performance of subjects across treatments (relative to the number of tasks they have seen) is similar.

6 Conclusion

In this paper we test what type of information people keep track of in problems with sequential learning of information. We show that individuals flexibly adjust what kinds of information they choose to remember. We show both theoretically and experimentally that decision-makers rationally adjust their memory strategy to the environment. In environments with relatively many signals and a clearly defined task people remember posteriors, while in settings with relatively few signals and uncertainty about the final decision people remember signals. This provides a foundation for understanding when different assumptions about bounded memory apply. In particular, we demonstrate that the two assumptions frequently made in the literature, that either people remember only posterior beliefs or they only remember individual signals, can both be rationalized in different environments.

References

Bays, P. M. and Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science*, 321(5890):851–854.

Bénabou, R. and Tirole, J. (2004). Willpower and personal rules. *Journal of Political Economy*, 112(4):848–886.

Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. In *Handbook of Behavioral Economics - Foundations and Applications 2*, pages 69–186. Elsevier.

Bergemann, D. and Valimaki, J. (2006). Bandit problems. *SSRN Electronic Journal*.

Bordalo, P., Conlon, J., Gennaioli, N., Kwon, S. Y., and Shleifer, A. (2023). Memory and probability. *The Quarterly Journal of Economics*, 138(1):265–311.

Brady, T. F., Konkle, T., and Alvarez, G. A. (2009). Compression in visual working memory: using statistical regularities to form more efficient memory representations. *Journal of Experimental Psychology: General*, 138(4):487.

Bénabou, R. and Tirole, J. (2002). Self-confidence and personal motivation. *The Quarterly Journal of Economics*.

Chatterjee, K., Guryev, K., and Hu, T.-W. (2022). Bounded memory in a changing world: Biases in behaviour and belief. *Journal of Economic Theory*, 206:105556.

Chatterjee, K. and Hu, T.-W. (2023). Learning with limited memory: Bayesianism vs heuristics. *Journal of Economic Theory*, 209:105642.

Chew, S. H., Huang, W., and Zhao, X. (2020). Motivated false memory. *Journal of Political Economy*, 128(10):3913–3939.

Cover, T. and Hellman, M. (1970). The two-armed-bandit problem with time-invariant finite memory. *IEEE Transactions on Information Theory*, 16(2):185–195.

Cowan, N. (2010). The magical mystery four: How is working memory capacity limited, and why? *Current directions in psychological science*, 19(1):51–57.

Da Silveira, R. A., Sung, Y., and Woodford, M. (2020). Optimally imprecise memory and biased forecasts. *NBER Working Paper*.

d'Acremont, M., Schultz, W., and Bossaerts, P. (2013). The human brain encodes event frequencies while forming subjective beliefs. *Journal of Neuroscience*, 33(26):10887–10897.

Deck, C., Jahedi, S., and Sheremeta, R. (2021). On the consistency of cognitive load. *European Economic Review*, 134:103695.

Dow, J. (1991). Search decisions with limited memory. *The Review of Economic Studies*, 58(1):1–14.

Enke, B., Schwerter, F., and Zimmermann, F. (2023). Associative memory, beliefs and market interactions. *Working Paper*.

Ericsson, K. A. and Kintsch, W. (1995). Long-term working memory. *Psychological review*, 102(2):211.

Fudenberg, D., Lanzani, G., and Strack, P. (2022). Selective memory equilibrium. *SSRN Electronic Journal*.

Gottlieb, D. (2014). Imperfect memory and choice under risk. *Games and Economic Behavior*, 85:127–158.

Graeber, T., Zimmermann, F., and Roth, C. (2022). Stories, statistics, and memory. *SSRN Electronic Journal*.

Hu, T.-W. (2023). Forgetful updating and stubborn decision-makers. *Economic Theory*, 75(3):781–802.

Kahana, M. J. (2012). *Foundations of Human Memory*. OUP USA.

Kocer, Y. (2010). Endogenous learning with bounded memory. *Economic Theory Center Working Paper*, (001-2011).

Leung, B. T. K. (2023). A simple model of memory-based beliefformation. *Working Paper*.

Ma, W. J., Husain, M., and Bays, P. M. (2014). Changing concepts of working memory. *Nature neuroscience*, 17(3):347–356.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review*, 63(2):81.

Monte, D. (2005). Reputation with a bounded memory receiver. *Working Paper*.

Monte, D. (2013). Bounded memory and permanent reputations. *Journal of Mathematical Economics*, 49(5):345–354.

Monte, D. (2014). Learning with bounded memory in games. *Games and Economic Behavior*, 87:204–223.

Monte, D. and Said, M. (2010). Learning in hidden markov models with bounded memory. *MPRA Paper*, (23854).

Mullainathan, S. (2002). A memory-based model of bounded rationality. *The Quarterly Journal of Economics*, 117(3):735–774.

Neligh, N. L. (2022). Rational memory with decay. *SSRN Electronic Journal*.

Oberauer, K., Farrell, S., Jarrold, C., and Lewandowsky, S. (2016). What limits working memory capacity? *Psychological bulletin*, 142(7):758.

Rothschild, M. (1974). Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy*, 82(4):689–711.

Shadlen, M. N. and Shohamy, D. (2016). Decision making and sequential sampling from memory. *Neuron*, 90(5):927–939.

Spaniol, J., Schain, C., and Bowen, H. J. (2014). Reward-enhanced memory in younger and older adults. *Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 69(5):730–740.

Wachter, J. A. and Kahana, M. J. (2023). Associative learning and representativeness. *SSRN Electronic Journal*.

Wald, A. (1945). Sequential method of sampling for deciding between two courses of action. *Journal of the American Statistical Association*, 40(231):277–306.

Wieth, M. and Burns, B. D. (2006). Incentives improve performance on both incremental and insight problem solving. *Quarterly Journal of Experimental Psychology*, 59(8):1378–1394.

Wilson, A. (2014). Bounded memory and biases in information processing. *Econometrica*, 82(6):2257–2294.

A Tables and Figures

A.1 Literature

Paper	T/E?	Constraint	Context
Assumption: Remember posterior belief(s)			
Cover and Hellman (1970)	t	finite memory states	2 states of the world, question known, potentially many signals
Dow (1991)	t	finite memory states	finite (ordered) states, question known, potentially many signals
Monte (2005)	t	finite memory states	2 states of the world, question known, potentially many signals
Kocer (2010)	t	finite memory states	2 states of the world, question known, potentially many signals, decision each period
Monte and Said (2010)	t	finite memory states	2 states of the world, question known, potentially many signals, decision each period
Monte (2013)	t	finite memory states	2 states of the world, question known, potentially many signals
Monte (2014)	t	finite memory states	2 states of the world, question known, potentially many signals
Wilson (2014)	t	finite memory states	2 states of the world, question known, potentially many signals
Benjamin (2019)	e	not modelled	2 states of the world, question known, varied other context
Chatterjee et al. (2022)	t	finite memory states	2 states of the world, question known, potentially many signals
Chatterjee and Hu (2023)	t	finite memory states	2 states of the world, question known, potentially many signals
Hu (2023)	t	finite memory states	2 states of the world, question known, potentially many signals
Assumption: Remember signals			
Mullainathan (2002)	t	selective recall based on similarity	question known ex-ante? potentially many signals

Paper	T/E?	Constraint	Context
Bénabou and Tirole (2004)	t	imperfect memory with potential manipulation	Finite ordered states of the world 1 signal motivated reasoning
Gottlieb (2014)	t	imperfect memory with potential manipulation	Finite ordered states of the world 1 signal motivated reasoning
Chew et al. (2020)	t	imperfect memory with potential manipulation	Finite ordered states of the world 1 signal motivated reasoning
Fudenberg et al. (2022)	t	selective recall	Finite states potentially many signals decision each period motivated reasoning
Graeber et al. (2022)	e	memory with decay	2 states of the world, 1 signal, 1 question known ex-ante, 3 tasks simultaneously
Neligh (2022)	t	memory with decay	Finite states, potentially many signals, decision each period
Bordalo et al. (2023)	t&e	selective recall based on similarity	question not known ex-ante, many signals (40)
Enke et al. (2023)	e	selective recall based on context similarity	between 0 and 3 signals, 1 question known ex-ante, ordered states, 12 tasks simultaneously
Leung (2023)	e	imperfect recall	2 (or more) states of the world, few/many signal, question not known ex-ante
Wachter and Kahana (2023)	t	selective recall based on similarity	finite states, question not known
Endogenous: Posterior beliefs or signals			
Da Silveira et al. (2020)	t	memory cost based on complexity	Finite ordered states of the world, question known, potentially many signals, decision each period

Table 1: Overview of literature on sequential information processing (with memory constraints). We include different papers in which a single decision maker receives several noisy signals about some state of the world over time. Papers are classified according to whether decision-makers are assumed to remember posteriors beliefs (i.e. summary statistics about states), signals, or can choose. The first column provides the paper, the second classifies the contribution as theoretical (t) or empirical (e), the third describes the memory constraint, and the final provides the additional context of the decision environment. For additional discussion please see Section 2.

A.2 Main Hypothesis

	Dependent variable:
	Question Choice: Posterior
Constant	2.324^{**}
	(1.022)
Treatment: Signals	-1.737^{***}
	(0.626)
Distraction: Math tasks solved	-0.046^{**}
	(0.023)
Distraction: Memory tasks solved	-0.016
	(0.070)
Duration	0.004
	(0.014)
Button clicked: Posterior question	0.816
	(0.508)
Button clicked: Posterior q. - scoring rule	1.353^{***}
	(0.464)
Button clicked: Signal question	-1.636^{***}
	(0.496)
Gender: Male	0.484
	(0.384)
Gender: Other	-0.050
	(1.036)
Age	0.004
	(0.018)
Use of Hints	0.014^{*}
	(0.007)
Remembering Every Hint	-0.010
	(0.007)
Remembering Easy	-0.376
	(0.305)
Updating Extent	0.001
	(0.006)
Updating Easy	-0.324
	(0.267)
Observations	100
Log Likelihood	-46.128
Akaike Inf. Crit.	124.256
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 2: Probit Regression Output

Variable Name	Marginal Effect	SE	P value
Treatment: Signals	-0.46	0.13	0.00
Distraction: Math tasks solved	-0.01	0.01	0.03
Distraction: Memory tasks solved	-0.00	0.02	0.82
Duration	0.00	0.00	0.77
Button clicked: Posterior question	0.19	0.10	0.06
Button clicked: Posterior q. - scoring rule	0.33	0.09	0.00
Button clicked: Signal question	-0.43	0.10	0.00
Gender: Male	0.12	0.09	0.19
Gender: Other	-0.01	0.27	0.96
Age	0.00	0.00	0.84
Use of Hints	0.00	0.00	0.05
Remembering Every Hint	-0.00	0.00	0.14
Remembering Easy	-0.10	0.08	0.21
Updating Extent	0.00	0.00	0.92
Updating Easy	-0.08	0.07	0.22

Table 3: Marginal effects from the regression output reported in Table 2

A.3 Additional Results

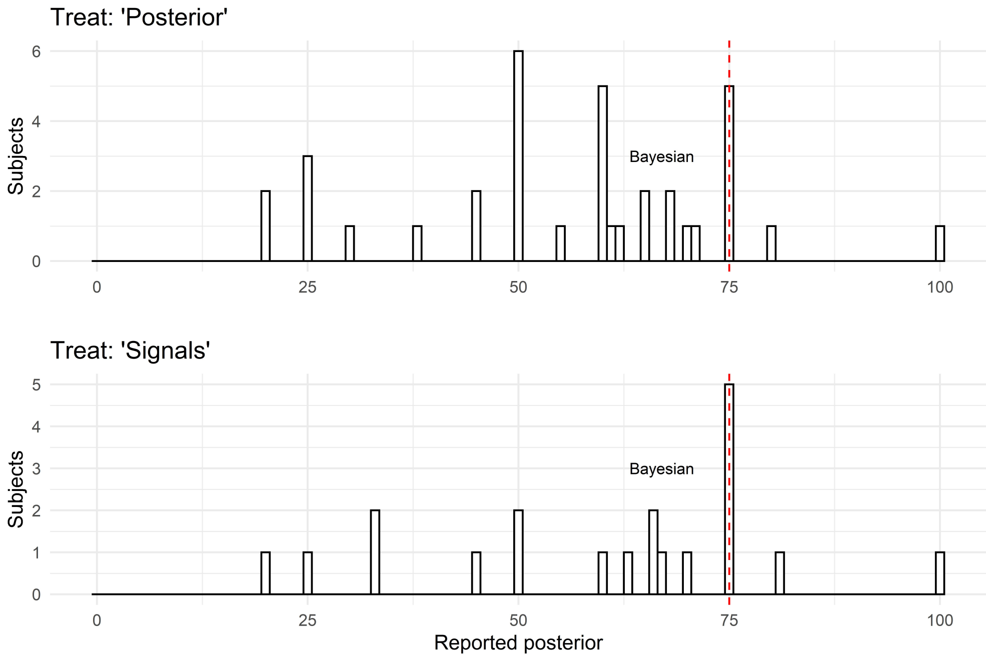

Figure 5: Histogram of reported posterior beliefs, split by treatment. The dashed line at 75% indicates the Bayesian posterior belief.

A.4 Distraction Task

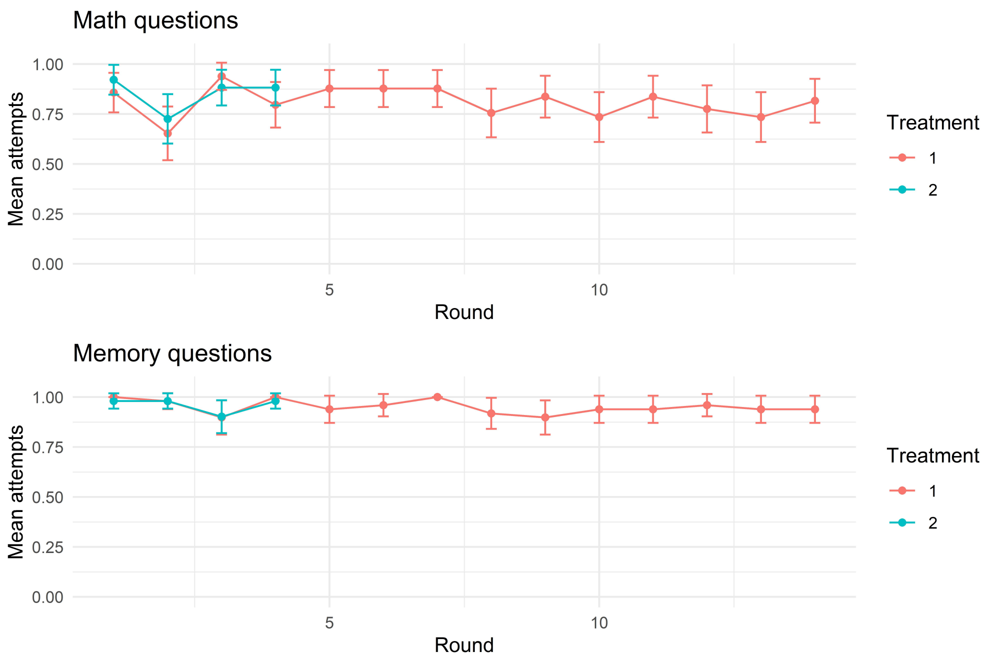

Figure 6: Mean attempts for the math and memory distraction tasks over time, split by treatment.

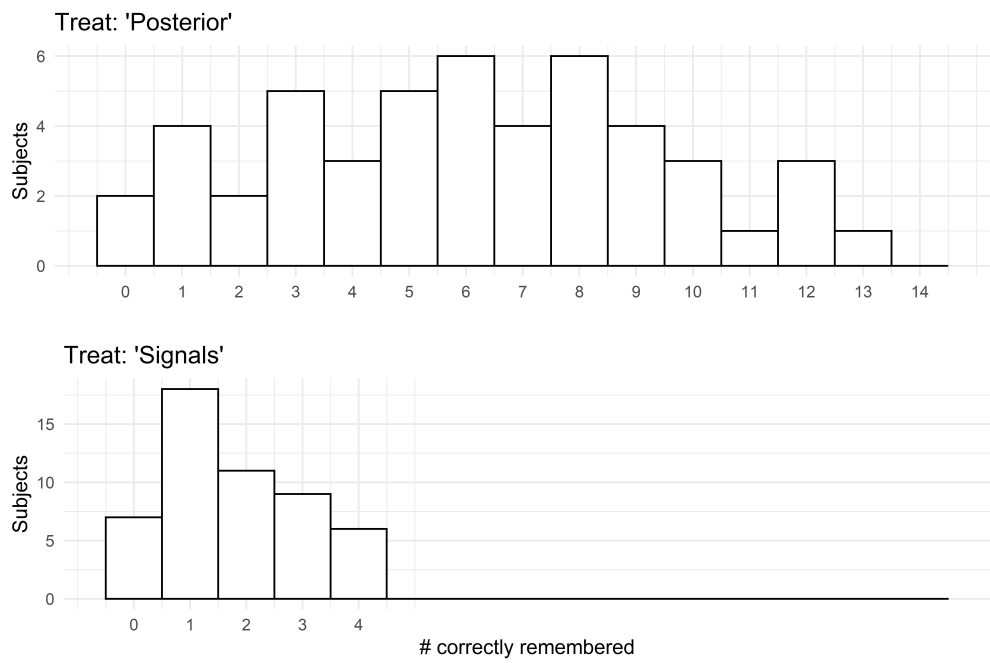

Figure 7: Histogram of correct answers in the memory distraction tasks over time, split by treatment. Note that the maximum number of correct answers is 14 for treatment 'posterior' and 4 for treatment 'signals'.

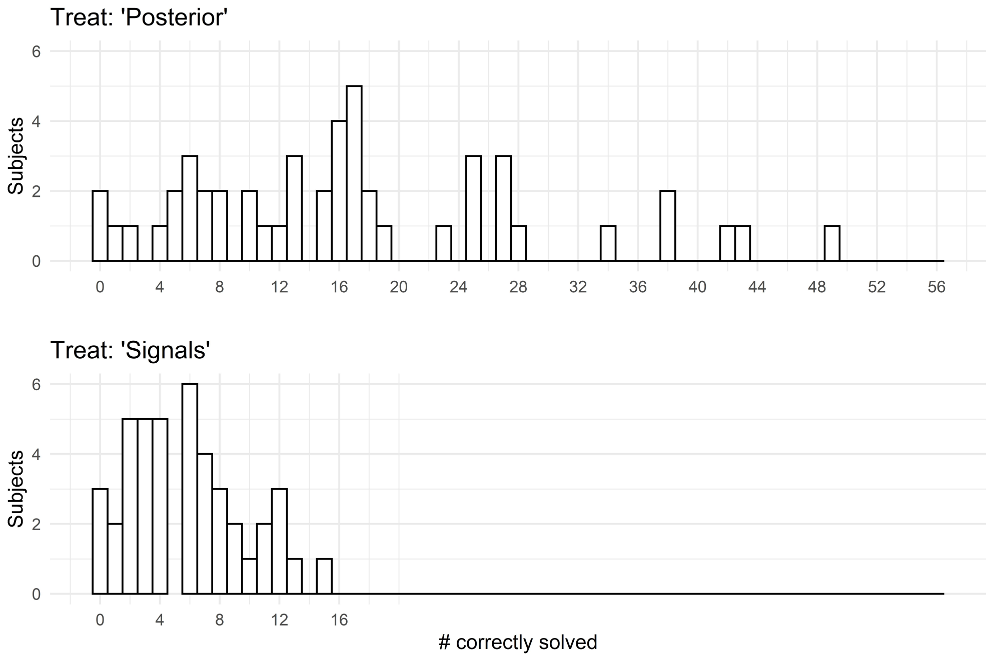

Figure 8: Histogram of correct answers in the math distraction tasks over time, split by treatment. Note that the maximum number of correct answers is 56 for treatment 'posterior' and 16 for treatment 'signals'.

B Instructions and Screenshots

B.1 Instructions - Treatment 'Posterior'

A group of people is tasked with selecting several winning numbers for a local raffle. The group is composed of members from three different teams. One person from each team is randomly selected (all with equal chance). Below are the three teams and their members:

Team names	**Members**
'Size'	Hugh & Loa
'Parity'	Eve & Todd
'Round'	Iris & Ron

All three selected people meet up and together select the winning numbers. The selected numbers will be shown to you.

Your task:

At the end, after seeing all numbers, you will be asked: **Based on the numbers you have seen, guess** which person from **Team 'Parity'** was randomly selected.

Hints:

Together the three selected people choose **14 different numbers between 1 and 100**. Each person likes some numbers more than others. Together they make sure that everyone likes the 14 chosen numbers. These 14 numbers will be shown to you.

Below you can find more information about the team members and the numbers they like to select:

Team 'Size'		Team 'Parity'		Team 'Round'	
Hugh	Loa	Eve	Todd	Iris	Ron
Hugh likes **high numbers** (51, ..., 100)	Loa likes **low numbers** (1, ..., 50)	Eve likes **even numbers** (2, 4, ..., 100)	Todd likes **odd numbers** (1, 3, ..., 99)	Iris likes **irregular numbers**, i.e. not divisible by 5 (1, 2, ..., 99)	Ron likes **round numbers**, i.e. divisible by 5 (5, 10, ..., 100)
On average, 75% of numbers selected by groups with Hugh are high, and 25% are low.	On average, 25% of numbers selected by groups with Loa are high, and 75% are low.	On average, 75% of numbers selected by groups with Eve are even and 25% are odd.	On average, 25% of numbers selected by groups with Todd are even and 75% are odd.	On average, 75% of numbers selected by groups with Iris are irregular and 25% are round.	On average, 25% of numbers selected by groups with Ron are irregular and 75% are round.

Example hint: 'The [first] winning number chosen by the three randomly selected people is: [23]'

Remember:

At the end, after seeing all numbers, you will be asked: **Based on the numbers you have seen, guess** which person from **Team 'Parity'** was randomly selected.

B.2 Instructions - Treatment 'Signals'

A group of people is tasked with selecting several winning numbers for a local raffle. The group is composed of members from three different teams. One person from each team is randomly selected (all with equal chance). Below are the three teams and their members:

Team names	**Members**
'Size'	Hugh & Loa
'Parity'	Eve & Todd
'Round'	Iris & Ron

All three selected people meet up and together select the winning numbers. The selected numbers will be shown to you.

Your task:

At the end, after seeing all numbers, you will be asked one of the three questions below.

Based on the numbers you have seen...

- ... **guess** which person from **Team 'Size'** was randomly selected, or
- ... **guess** which person from **Team 'Parity'** was randomly selected, or
- ... **guess** which person from **Team 'Round'** was randomly selected.

All three questions are equally likely to be selected.

Hints:

Together the three selected people choose **14 different numbers between 1 and 100**. Each person likes some numbers more than others. Together they make sure that everyone likes the 14 chosen numbers. These 14 numbers will be shown to you.

Below you can find more information about the team members and the numbers they like to select:

Team 'Size'		Team 'Parity'		Team 'Round'	
Hugh	Loa	Eve	Todd	Iris	Ron
Hugh likes **high numbers** (51, ..., 100)	Loa likes **low numbers** (1, ..., 50)	Eve likes **even numbers** (2, 4, ..., 100)	Todd likes **odd numbers** (1, 3, ..., 99)	Iris likes **irregular numbers**, i.e. not divisible by 5 (1, 2, ..., 99)	Ron likes **round numbers**, i.e. divisible by 5 (5, 10, ..., 100)
On average, 75% of numbers selected by groups with Hugh are high, and 25% are low.	On average, 25% of numbers selected by groups with Loa are high, and 75% are low.	On average, 75% of numbers selected by groups with Eve are even and 25% are odd.	On average, 25% of numbers selected by groups with Todd are even and 75% are odd.	On average, 75% of numbers selected by groups with Iris are irregular and 25% are round.	On average, 25% of numbers selected by groups with Ron are irregular and 75% are round.

Example hint: 'The [first] winning number chosen by the three randomly selected people is: [23]'

Remember:
At the end, after seeing all numbers, you will be asked one of the three questions below.
Based on the numbers you have seen...

- ... **guess** which person from **Team 'Size'** was randomly selected, or
- ... **guess** which person from **Team 'Parity'** was randomly selected, or
- ... **guess** which person from **Team 'Round'** was randomly selected.

B.3 Instructions - Comprehension Questions (both treatments)

To make sure you correctly understood the instructions for the main task, please answer the following 4 questions. You can read the instructions once more by clicking on the button below:
[Button: 'Show/hide instructions']

If you answer 2 or more questions incorrectly you will not be eligible for a bonus payment in this study.

You will only be able to proceed once you answer all questions correctly. Otherwise you will be shown the instructions once again.

Which people are selected to decide the winning numbers for the local raffle?

- 3 out of the 6 are randomly selected.

- 3 are randomly selected, one from each team.
- The winning numbers are chosen randomly, the people do not matter.

Suppose the three selected people are: Hugh, Todd and Ron. Which types of numbers are more likely to be winning numbers? Numbers that are:

- Low, odd and irregular.
- High, even and round.
- Low, even and round.
- High, odd and round.

In the end I will be asked to guess...

- ...which person from Team 'Size' was selected.
- ...which person from Team 'Parity' was selected.
- ...which person from Team 'Round' was selected.
- One of the 3 questions above will be randomly selected.

How many winning numbers will be shown to me?

- 4
- 9
- 14

——————————— [Page break] ———————————

The correct answers are:

1. 3 are randomly selected, one from each team.
2. High, odd and round.
3. [T1:] ...which person from Team 'Parity' was selected. / [T2:] One of the 3 questions above will be randomly selected.
4. [T1:] 14 / [T2:] 4

B.4 Signals

The complete list of signals for treatment 1: 12, 7, 84, 37, 62, 75, 60, 8, 34, 9, 35, 88, 21, 30.
The complete list of signals for treatment 2: 12, 7, 84, 37.

Figure 9: Example signal as shown to subjects.

B.5 Question choice

Chance for additional payoff - your choice:

Before the experiment continues you have the chance to earn some additional money. You can choose which type of question you prefer to answer. You have the choice between two options:

- A question regarding the chance which person from Team 'Parity' was selected in the beginning:

——— [Button: 'Details'] ———

You will be asked to state the chance (in %) with which you think Eve/Todd was selected from Team 'Parity'. Your bonus payment (either \$8.00 or \$0.00) depends on the accuracy of your answer. The more accurate your answer is, the higher your expected payoff. Example: 'Based on the hints you received so far, what is the chance (in %) that Eve was randomly selected from Team 'Parity' in the beginning?'

——— [Button: 'Formula'] ———

You will receive the bonus payment of \$8.00 with some probability that depends on your report. The probability is calculated according to the following formula: $1 - (\text{error})^2$, where the error is the difference between your reported percentage and 100% if Eve was selected and 0% if Todd was selected.

You can maximize your expected earnings by reporting what you truly think the chances are that Eve was selected from Team 'Parity' in the beginning.

——— [End button: 'Formula'] ———

——— [End button: 'Details'] ———

- A question regarding one of the [T1:] 11 / [T2:] 3 numbers you have previously seen: [Button: 'Details']

——— [Button: 'Details'] ———

You will be asked to recall one of the [T1:] 11 / [T2:] 3 winning numbers that was shown to you so far. Example: 'What was the [third] winning number selected by the group of people?' One of the numbers you saw so far is randomly selected for this question. If your answer is correct you will receive \$8.00, otherwise you will receive \$0.00.

——— [End button: 'Details'] ———

Please select one of the two options below. This will be the type of question asked to you on the next page. With either type of question you can earn an additional payoff of \$8.00 depending on your answer to the selected question.

- A question regarding the chance which person from Team 'Parity' was selected in the beginning.
- A question regarding one of the numbers I have previously seen.

[Question regarding chance:]

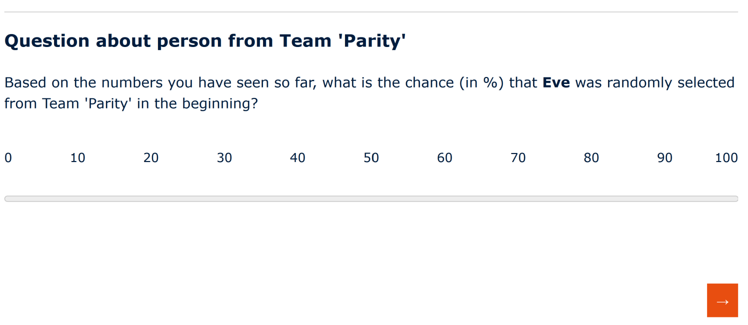

Figure 10: Posterior question

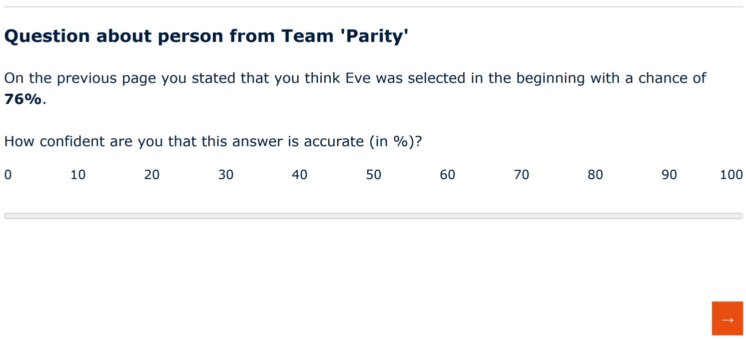

Figure 11: Confidence in reported posterior (not incentivized)

[Question regarding one of the numbers:]

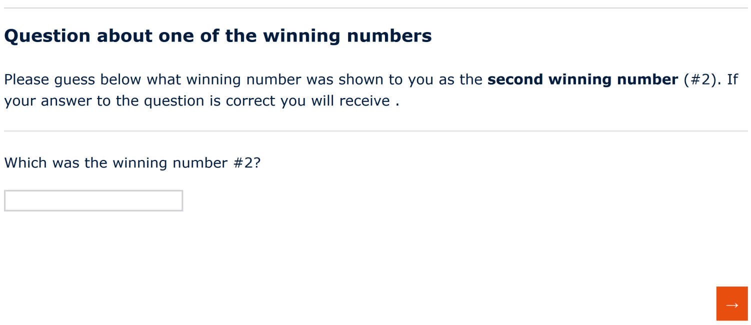

Figure 12: Elicitation of signal memory

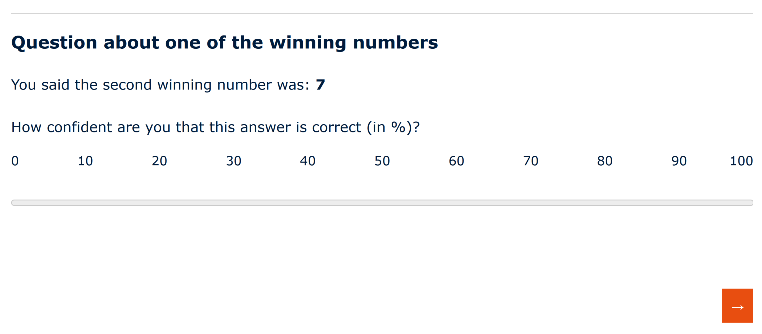

Question about one of the winning numbers

You said the second winning number was: **7**

How confident are you that this answer is correct (in %)?

0 10 20 30 40 50 60 70 80 90 100

→

Figure 13: Confidence in memory of signal (not incentivized)

B.6 Distraction Task

The distraction task is shown to subjects immediately after each signal. It includes a number memory task and four simple calculation questions. Each time subjects are shown the instructions. The following four pages are shown to subjects in sequence.

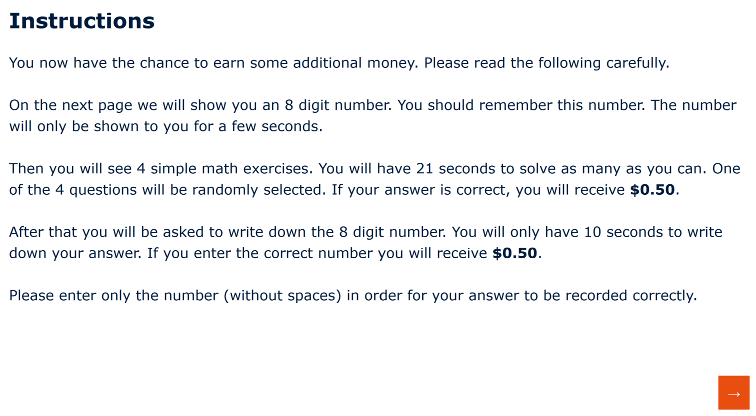

Instructions

You now have the chance to earn some additional money. Please read the following carefully.

On the next page we will show you an 8 digit number. You should remember this number. The number will only be shown to you for a few seconds.

Then you will see 4 simple math exercises. You will have 21 seconds to solve as many as you can. One of the 4 questions will be randomly selected. If your answer is correct, you will receive **$0.50**.

After that you will be asked to write down the 8 digit number. You will only have 10 seconds to write down your answer. If you enter the correct number you will receive **$0.50**.

Please enter only the number (without spaces) in order for your answer to be recorded correctly.

→

Figure 14: Instructions for distraction task

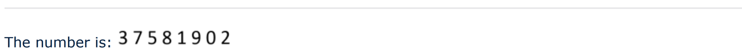

The number is: 3 7 5 8 1 9 0 2

Figure 15: First part of memory distraction task

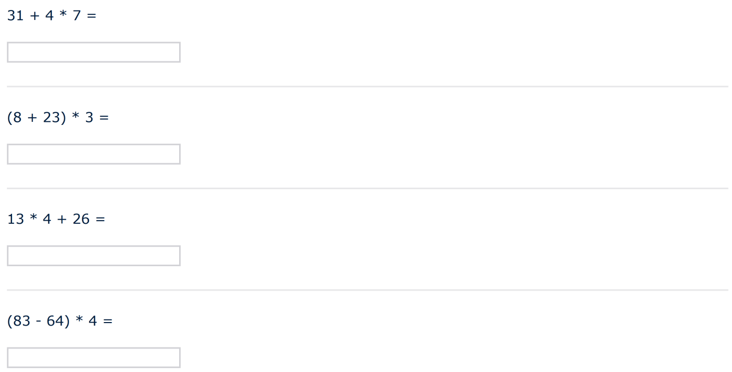

Figure 16: Math distraction task

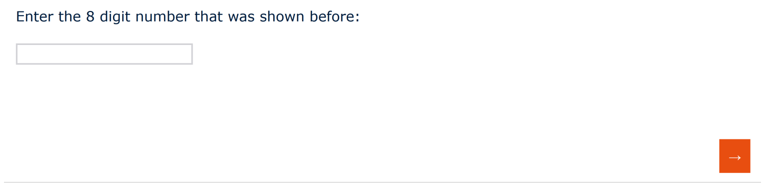

Figure 17: Second part of memory distraction task

B.7 Post-experiment survey

[First page]

What is your gender?

Options: Female; Male; Other

What is your age? Please enter the number below.

What country were you born in? Please select the country name below.

What is the main field of study for your undergraduate degree?

Options: Management/Business; Economics; Humanities; Liberal Arts; Education; Engineering; Science; Social Science; Agriculture; Pharmacy; Nursing; Other

---

What is your GPA?

Options: 3.5-4.0; 3.0-3.5; 2.5-3.0; 2.0-2.5; Below 2.0

Are you an undergraduate student (which year) or a graduate student?

Options: First year; Second year; Third year; Fourth year or above; Graduate student

---

What is your SAT score? Please enter the number as best as your remember below.

[Second Page]

On this page you will be asked a few questions regarding your behavior and strategy in this survey.

You saw several numbers that provided information about which people were selected in the beginning. To what extent were you trying to use the information contained in these hints?

0: I did not use the information at all.

100: I used all the information contained in the hints.

[slider]

To what extent were you trying to remember every number shown to you?

0: I did not try to remember each individual hint.

100: I tried to perfectly keep track of all hints and the dates.

[slider]

Did you think it was difficult to keep track of the numbers?

Options: Very difficult; Somewhat difficult; Somewhat easy; Very easy

---

After seeing each number, to what extent were you trying to keep track of which people were selected in the beginning?

0: I did not keep track at all of which people I thought were selected.

100: I tried to keep track precisely of which people I thought were the selected one(s).

[slider]

Did you think it was difficult to keep track of which people were the selected one(s)?

Options: Very difficult; Somewhat difficult; Somewhat easy; Very easy

---

You could choose between two questions to earn an additional payoff after several numbers were shown to you. You chose for: 'XX'. Why did you choose this question? Please explain your thought process as clearly as possible.

Do you have any general feedback regarding this survey? Was there anything that you did not fully understand or should be improved in the future?

[Final page]

Payoff Summary

Show up fee:

Payoff from math questions:

Payoff from number memory questions:

Payoff from main game:

Total:

Please record this number on your payment sheet and come to the front.