1 Experimental Details

We conduct an experiment in which decision makers engage in a repeated memory task (remembering lists of English nouns). We first describe the overall paradigm in subsection 1.1, and then the details of the experimental task in subsection 1.2

1.1 Experimental Paradigm

The timing of each round of the task is divided into three phases: 1) stimulus, 2) delay, 3) response. Thus, in a typical round, participants are first presented with a stimulus (a list of words) and invited to study and memorize it for a period of time. Then there is a delay (during which a difrerent task will be presented to the participants); finally, participants are asked to identify which of the words on the simulus list are present on a new word list, composed of both words from the stimulus list, and new words. At this point the round end, and a new one (with a new word list being presented) begins.

During the stimulus phase, the information is available to the individuals. In this phase information gathering and some amount of memory encoding takes place. During the delay phase, the initial stimulus is removed. In this phase, the remainder of the encoding and memory maintenance takes place. During the response phase, participants are required to complete the task. Memory retrieval and action take place during this phase.

In the experiment, $r \in \{L, H, H'\} = \{0, 1, 2\}$ is the reward for successfully completing the task, and in any single block of rounds it only takes on two values. The payoff for a correct response can be low (0) or high (either H or H', but not both) with some probability. The player is not informed of which level r takes until some point in the experiment determined randomly. However, *only low rewards are ever revealed.* If rewards are high, this fact will never be revealed. If rewards are low then the following will occur each one quarter of the time: (1) Low reward is revealed before stimulus. (2) Low reward is revealed after stimulus and before delay. (3) Low reward is revealed after the delay and before the response. (4) Reward is never revealed.

There are two blocks of 20 rounds of this scheme; the different between the blocks is the level of the high reward. In one block it is H and in the other it is H'.

In addition to the word recall questions, there are two other things that can happen. Half of the time, instead of the "typical" round described above, the response will be requested immediately after the stimulus is removed, with no delay period (in this case reward timing 2 and 3 will behave the same). In addition, half the time, after a response is given, the same response will be requested a second time, this time with no reward. Note, the probability of these events are independent of one another and independent of reward timing and reward level.

1.2 Experimental Task

The experimental task is a verbal recall task that proceeds in rounds. We use the Toronto Word Pool as a source of various commonly used words to construct a "grand" word list of 50 English nouns. In a treatment we first present the participants with the entire "grand" word list. This feature allows us to control the state space, because we provide the participants with all possible words they might see in advance; without this feature the state space is potentially very large (the set of all nouns in American English, or some subset of it that may be salient to the participant). In each round a specific subset (10 words) of the "grand" word list will be selected (each word has an equal probability of being selected to be in the "round" word list; the selections are made without replacement within a list, but independently across the lists), and shown to the participant during the "stimulus" phase of the round.

The subset lists are pre-randomized (by setting a seed during the selection process), fixed, and thus are the same across participants. This allows us to compare performance on comparable tasks across individual participants.

The recall phase resembles a multiple-price list setup: we present participants with a second list (of ten words), in which the words are drawn half the time from the sublist for that round, and the other half of the time from the rest of the grand list. This means that, the conditional probability - the belief - of whether a word is on the sublist for a particular round is $\frac{1}{2}$.

The true state of the world that is relevant to a given memory task is the sublist drawn for that memory task. In our modeling and analysis we make the simplifying assumption that each word has an independent (one in fifty) chance of being on the list (ignoring the fixed length). This assumptions dramatically simplifies notation and analysis and should be innocuous.

With this assumption we can essentially treat every memory task as fifty parallel inference tasks with a joint state space of $\{0, 1\}^{50}$, since every word is either in the list or not in the list. Of these fifty potential inference tasks, ten are checked with each response list.

In between the stimulus and the recall phases there is a delay phase. During the delay phase participants are asked to attend to a "distractor task." During this period, there is a blue square on the screen, which periodically turns green for three seconds. The instance when this happens is random (drawn uniformly over the length of the delay period). The participants are asked to press the "space" bar on the computer in front of them whenever this happens; if they press the space bar when the screen is green (i.e. during a small time period), they receive an extra $0.50 which gets

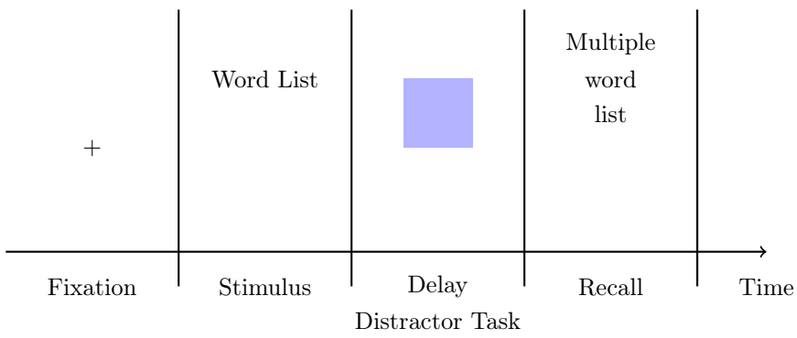

Figure 1: Experimental Paradigm: Overview of a Typical Round

added to their total earnings from the experiment. The timing of the instances when the square changes color is independent and identically distributed. The delay period is important in creating conditions under which effort choices are identified:

The reason for including the remunerated distractor task during the delay period is three-fold. First, this allows us to ensure that participants are not using working memory to store the word lists - the task distracts them from it, emphasizing and relying on visual attention rather than verbal processing. Second, this allows us to view the results of the recall component as comparable, because all participants were engaged in the same, uniform, known, attention-consuming task during the delay period. The fact that the distractor task is also remunerated and incentivized ensures that participants complete the task.

Participants are reimbursed for their decisions via a "probability points" scheme. There is a $10 show-up fee, and a $50 prize. Each correct answer (on both the distractor tasks) and a correct answer on a word task for which the reward was high increases the probability of getting the prize by 0.05. Thus, at the end of the experiment each participant will get the show-up fee, plus a lottery in which they get the $50 prize with a probability equal to the number of correct answers on the questions for which the reward is high plus the number of correct answers on the distractor task times 0.05.