# Learning under Ambiguity when Information Acquisition is Costly: an Experiment

Konstantinos Georgalos* John Hey†

October 16, 2019

## 1 Introduction

In an extension of the standard Ellsberg thought experiment, we experimentally test the case where it is possible for the decision maker to obtain some information regarding the composition of the ambiguous urn, before placing her bets on the an urn and a colour. More particularly, the DM can postpone their decisions in order to observe realizations from a diffusion process whose drift is equal to the proportion of red balls in the unknown (ambiguous) Bag. Epstein and Ji (2019) develop a new theoretical model in which they show that subject to the degree of a decision maker's ambiguity aversion, it can be optimal to reject learning completely, and, if some learning is optimal, then it is never optimal to bet on the risky Bag after stopping. TBC

*Department of Economics, Lancaster University Management School, Lancaster, LA1 4YX, U.K.
✉ k.georgalos@lancaster.ac.uk, ☎ +44 (0)15245 93170
†Department of Economics, University of York, Heslington, York, YO10 5DD, U.K.
✉ john.hey@york.ac.uk, ☎ +44 (0) 01904 323786

## 2 Related Literature

Trautmann and Zeckhauser (2013)

Ert and Trautmann (2014)

Baillon et al. (2018)

Nicholls et al. (2015)

Engle-Warnick and Laszlo (2017)

Fudenberg et al. (2018)

## 3 Theoretical Framework

The decision maker (DM) faces two Bags containing red and blue balls, a risky one in which the proportion of red balls is $1/2$ and an ambiguous one in which the colour composition is unknown. The unknown proportion of red balls is $1/2 + \theta$ with $\theta \in [-1/2, 1/2]$. $\theta$ represents the bias towards red, therefore $\theta > 0$ indicates more red than blue, while $\theta < 0$ the opposite. When $\theta = 0$ the composition of the Bag is identical to the known one. There is ambiguity about $\theta$ modeled by the set of priors $\mathcal{M}_0$. The number of balls is assumed to be large in order to treat $\theta$ as a continuous variable.

The DM has to choose whether to bet on the risky or the ambiguous Bag and then choose the colour on which she is going to bet on. She can postpone the choice so that she can learn about $\theta$ by observing realisations $Z$ of a signal process generated by a standard Brownian motion:

$$Z_t = \theta_t + \sigma B_t \tag{1}$$

There is a constant per-unit-time cost $c > 0$ of learning. Her choice of when to stop is described by a stopping time (or strategy) $\tau$. If the DM stops learning at $t$, her conditional expected payoff is $X_t$ and represents the indirect utility she can attain by choosing optimally between the available bets at $t$ (bet on red or blue, both from the ambiguous Bag).

The DM is a maxmin agent who is forward looking and applies backward induction, and is solving the following optimal stopping problem:

$$\max_{\tau} \min_{P \in \mathcal{P}_0} E_P(X_\tau - c\tau) \tag{2}$$

assuming that bets have prizes 0 and 1, and the utility index $u(.)$ is normalised such that $u(0) = 0; u(1) = 1$.

At time $t$, the utility of betting on red from the ambiguous Bag is $\min_{\mu \in \mathcal{M}_t} E_\mu$ where

$$E_\mu \equiv \int (1/2 + \theta) d\mu \tag{3}$$

similarly $E^*_\mu \equiv \int (1/2 - \theta) d\mu$ for betting on blue.

In order to obtain closed form solutions, Epstein and Ji assume that the Bag is biased and the proportion of red is either $\delta_{-\alpha} = 1/2 - \alpha$ or $\delta_\alpha = 1/2 + \alpha$, with $0 < \alpha < 1/2$ but there is ambiguity $\epsilon$ about which direction ($\epsilon \in (0, 1)$).

The set of prior beliefs is given by:

$$\mathcal{M}_0 = \{(1 - m)\delta_{-\alpha} + m\delta_\alpha : \frac{1 - \epsilon}{2} \leq m \leq \frac{1 + \epsilon}{2}\} \tag{4}$$

$\epsilon$ represents ambiguity aversion and the set $\mathcal{M}_0$ can be identified with the probability interval $[\frac{1-\epsilon}{2}, \frac{1+\epsilon}{2}]$. When $\epsilon = 0$ the DM is a Bayesian who faces uncertainty with variance $\alpha^2$ about the true bias, but no ambiguity. $\alpha$ measures the degree of this prior uncertainty.

They then provide the formula of how the set of priors is updated by Bayesian updating and then solve the optimal stopping problem. By defining $l(r)$ as:

$$l(r) = 2\log(\frac{r}{1 - r}) - \frac{1}{r} + -\frac{1}{1 - r} \tag{5}$$

with $r \in (0,1)$ and $\hat{r}$ as:

$$l(\hat{r}) = -\frac{2\alpha^3}{c\sigma^2}$$

they obtain the following theorem:

1. $\tau^* = 0$ if and only if $\frac{1+\epsilon}{2} \geq \hat{r}$.

2. Let $\frac{1+\epsilon}{2} \leq \hat{r}$ then the optimal stopping rule satisfies $\tau^* > 0$ which is given by $\tau^* = \min\{t \geq 0 : |Z_t| \geq \bar{z}\}$ with:

$$\bar{z} = \frac{\sigma^2}{2\alpha}[log(\frac{1+\epsilon}{1-\epsilon}) + log(\frac{\bar{r}}{1-\bar{r}})] \tag{6}$$

   and $\bar{r}, \hat{r} < \bar{r} < 1$ the unique solution[1] of

$$l(r) + l(\frac{1+\epsilon}{2}) = \frac{4\alpha^3}{c\sigma^2} \tag{7}$$

   After stopping, either the bet on red is chosen (if $Z_{\tau^*} \geq \bar{z}$) or the bet on blue is chosen ($Z_{\tau^*} \leq \bar{z}$). It is never optimal to bet on the risky Bag at $\tau^* > 0$. Extreme levels of ambiguity aversion predict that people will never search, while Bayesians will search for some positive $\tau$.

There is a also a corollary (B1, Appendix B, page 21) stating that the DM stops sampling later in each of the following cases:

1. $c$ falls

2. $\epsilon$ increases in the interval $[0, 2\hat{r} - 1]$

3. $\sigma$ and $\alpha$ both increase in such away that $\frac{\alpha}{\sigma^2}$ is constant

[1] We use numerical methods to calculate $\bar{z}$.

## 4 Treatments and Experimental Procedures

To test the model, we conducted a lab experiment with 114 subjects (66 female) randomly allocated to four different treatments (we provide details below). Subjects were mostly undergraduate students from various fields of studies. Their average age was 22.1 years old. The experiment was computer-based, conducted in the Centre for EXEC laboratory at the University of York, with subjects recruited via the hroot recruiting system (Bock et al. 2014). Written Instructions (Appendix) were placed on the subjects' desks and these were read aloud by one of the experimenters over the tannoy system. Any questions were publicly answered; there were few.

The task in the experiment was identical to the one presented in section 3. In the front of the lab there were two bags, Bag A (the known bag) filled with 50 Red and 50 Blue plastic tokens and Bag B (the unknown Bag) filled with $50 + \alpha$ Red tokens and the remaining Blue. In practice there were three bags, as the subjects were aware that there are two versions of Bag B, one with $\alpha > 0$ (more Red than Blue tokens) and one with $\alpha < 0$ (more Blue than Red tokens). There were in total 20 rounds. In each round, a subject should choose a bag (either the known or the unknown) and a colour (either Red or Blue). Before the beginning of each round, the computer would randomly choose the sign of $\alpha$ (positive or negative). In each round the subjects could also buy information on whether the value of $\alpha$ is positive or negative, before choosing a bag and a colour. This information was provided in the form of a standard Brownian motion, as the one in Equation 1, where the subjects could observe realisations $Z$ of a signal process, at a cost $c$ for every second of sampling[2]. This signal process could

[2]In order to implement a Brownian motion in the lab we need to resort to a discrete time binomial approximation of the continuous stochastic process. Following Oprea et al (2009) the binomial approximation involves a fixed time interval $\Delta_t$ (tick) for each discrete step and two parameters corresponding to those of the Brownian process[3], namely the step size $h > 0$ of the proportional change in value and

be observed for a maximum of 60 seconds[4]. Subjects were provided with various examples of how the noisy signal would look like for different values of $\sigma$ (different levels of noise). Prior to the experiment, the subjects could participate in 5 practice rounds, in order to familiarise themselves with the task and the experimental software.

Subjects were paid based on a random incentive system. At the end of the experiment, one of the subjects would draw a ticket from a set of tickets numbered from 1 to 20, and this would determine the problem that the subjects would play out for real. The software would then recall which Bag and Colour each subject chose in that particular problem, and whether they spend money to obtain information. If the subject had chosen the unknown bag, the actual value of $\alpha$ would be revealed and the subject would be presented with the corresponding bag. The subject would then draw a token from his/her chosen bag on that problem (known or unknown) and if the colour of the token would be the same with the one he/she chose, the subject would be paid £$10 - tc$, otherwise $-tc$, with $t$ being the number of sampling seconds. Each subject received a show-up fee of £5.

Theorem 3.2 predicts that if the subject's ambiguity parameter is less than a particular upper bound, the subject will engage in sampling for some positive time $t$ until the diffusion process reaches a particular threshold. In addition, the sampling stops later when: (i) $c$ falls, and; (ii) $\alpha$ and $\sigma$ both increase in such away that $\frac{\alpha}{\sigma^2}$ is constant. To test the predictions of the theorem we vary the levels of $\alpha$ and $c$ in a $2 \times 2$ between-subjects design. More particularly, we implemented two levels of cost $c$, a Low-level

the uptick probability (the probability that the next tick will go up or down). $\Delta_t$ is set equal to 0.003 minutes which corresponds to 5 ticks per second. In our case, we do not specify the step size $h$ or $p$ as in the case of Oprea et al. since these are driven by the drift of the Brownian motion. Currently the version I have programmed generates a Brownian motion with $\Delta_t = 0.003$, $\mu$ equal to the bias of the red balls and $\sigma$ a parameter that we need to specify.

[4]The subjects were instructed that if the time is out during the sampling process, it will be assumed that they do not want to choose a bag and a colour in that particular problem. This happened in five separate cases.

cost of 2 pence per second, and a High-level cost of 4 pence per second. Similarly, we implemented two levels of $\alpha$, a Low-level where $\alpha$ could take the value $\pm 15$ (i.e. 65 Red tokens and 35 Blue tokens when $\alpha = 15$ or 35 Red tokens and 65 Blue tokens when $\alpha = -15$), and a High-level where $\alpha$ could take the value $\pm 30$. The value of $\sigma$ was adapted accordingly so that the ratio $\alpha/\sigma^2$ remained constant. There are in total 4 treatments LL, LH, HL and HH (the first element indicates the level of cost while the second one the level of $\alpha$). The details of the treatments are summarised in Table 1.

Based on the theoretical model and our experimental treatments, we summarise our testable hypotheses below. Given the extensive empirical evidence regarding the existence of non-neutral ambiguity attitudes, we expect that subjects will behave in a heterogeneous way, deviating from the limiting hypothesis of Bayesian DMs ($\epsilon = 0$). First, DMs with ambiguity neutral $\epsilon = 0$ or ambiguity averse attitudes ($0 < \epsilon < 2\hat{r} - 1$) will find it optimal to engage in sampling for some time, stop, and then bet on the ambiguous bag.

**Hypothesis 1.** *For intermediate values of $\epsilon$, it is optimal to sample as long as $Z_t \in [-\bar{z}, \bar{z}]$. When $Z_t$ hits either of the bounds, learning stops and the DM bets on the ambiguous bag.*

Then, DMs with extreme levels of ambiguity aversion ($\epsilon \geq 2\hat{r} - 1$), will never try to obtain information and will instead, bet directly on the risky bag.

**Hypothesis 2.** *If ambiguity (measured by $\epsilon$) is large enough relative to the payoffs , then no sampling is optimal and the DM bets on the risky bag immediately.*

The model makes the prediction that a DM who engaged in sampling, has obtained sufficient information, regarding the composition of the ambiguous bag, such that it is never optimal to stop and bet on the risky one.

**Hypothesis 3.** *A DM who spent money on sampling and obtained information, will never bet on the risky bag.*

Along with the ambiguity parameter $\epsilon$, three other parameters may have an effect on the size of sampling, the cost of sampling $c$, the bias $\alpha$ of Red balls and the noise of the signal $\sigma$. The next hypothesis states that a DM will stop sampling later, if the cost of sampling decreases.

**Hypothesis 4.** *As the cost of sampling $c$ decreases, the DM stops sampling later: $|Z_t^{LL}| > |Z_t^{HL}|$; and $|Z_t^{LH}| > |Z_t^{HH}|$.*

Finally, a higher level of both the bias $\alpha$ and the noise $\sigma$, will increase the size of sampling.

**Hypothesis 5.** *As $\alpha$ and $\sigma$ both increase in such a way that $\alpha/\sigma^2$ is constant, the DM stops sampling later: $|Z_t^{LH}| > |Z_t^{LL}|$; and $|Z_t^{HH}| > |Z_t^{HL}|$.*

# 5 Results

We report the results in two subsections. First we report some descriptive statistics and comparisons of the various treatments and then we proceed to a more formal test of the model, using structural econometric modelling techniques.

## 5.1 Descriptive statistics

Overall, 45 (39%) subjects sampled in all 20 repetitions, 11 (10%) subjects never sampled, 73 (64%) subjects sampled for 10 repetitions or more, 58 (51%) for 15 repetitions or more, while 25 (13%) subjects sampled for less than 5 repetitions.

Table 2 reports the number of trials in which subjects engaged in learning across the four treatments. Overall, in 65% of the trials, subjects sampled before choosing a bag. The highest percentage is observed in treatment HH (71%) while the lowest in LH (53%). If all subjects were behaving in a Bayesian way, then the percentage of sampling should have been close to 100%. The large variation across treatments provides evidence of non-neutral ambiguity attitudes across the experimental population, and therefore evidence in favour of the model's prediction (Hypothesis 1).

Table 3 reports the percentage of trials in which the subjects chose not to sample and then chose to bet on the risky bag. Overall, in 74% of the trials without sampling, the subjects found that no learning is optimal and they bet on the risky bag immediately. According to the model's definition, these subjects could be classified as extreme ambiguity averse (the value of $\epsilon$ exceeds the minimum threshold). Since there is a large number of repetitions, it is expected that there will be some noise in the data (subjects may choose to sample by mistake or to experiment in some rounds with sampling). To take this into consideration, we adopt a more flexible measure to classify these subjects as extreme ambiguity averse. More particularly, we count the number of subjects who sampled for less than 5 repetitions, and chose the risky bag for more than 75% of the repetitions without sampling. Based on this measure, 14 subjects (12%) can be classified as extreme ambiguity averse (Hypothesis 2). Remarkably in the treatment LL, in 48% of the trials the subjects did not sample and chose immediately the ambiguous bag, behaviour that could be classified as ambiguity seeking.

Table 4 reports the percentage of repetitions in which the subjects sampled for some time and then chose the risky bag. Overall, this happened in 14% of all the repetitions. (Hypothesis 3). Epstein and Ji (2019) argue that, if the bias is small, and if $Z_t$ is small over a large period $t$, then there is little to be gained by continuing sampling and the DM is at the same situation as at time 0. Therefore, if it is optimal to stop at $t$ and bet on the risky bag, it should have been optimal to do so at time 0 as well. Nevertheless, due to the specification of the set of priors $\mathcal{M}_0$ (the DM knows that the bias is $\pm\alpha$), they claim that it is never optimal to try learning for a while, stop and bet on the risky bag, otherwise the DM behaves in a dynamically-inconsistent way. A potential explanation of why subjects switched to the risky bag, is that despite the information they obtained, the signals were sufficiently close to 0, and were perceived as noise, rather than useful information regarding the direction of the bias. Indeed, notice in the last column of Table 4, in the treatments with low bias/volatility (LL and HL), the percentage of repetitions in which the subject samples and choses the risky bag is more than double compared to the treatments with high bias/volatility.

We now turn to Hypotheses 4 and 5. Table 5 reports the mean, trimmed mean at 10% level[5], the median and the standard deviation of the threshold (absolute value), while Table 6 reports the same information but in terms of total sampling time (in seconds). The model predicts that when the cost $c$ increases, the lengths of sampling decreases: $|Z_t^{LL}| > |Z_t^{HL}|$; and $|Z_t^{LH}| > |Z_t^{HH}|$. Between the treatments with low bias/volatility (LL,HL), treatment LL has the minimum mean threshold which therefore rejects the hypothesis $|Z_t^{LL}| > |Z_t^{HL}|$ ($p < 0.000$; Wilcoxon Rank-Sum test). In the high bias/volatility treatments (LH, HH), it holds that $|Z_t^{LH}| > |Z_t^{HH}|$ (0.130 vs. 0.110) and this difference is significant (p=0.065). Regarding Hypotheses 5, the theory predicts that higher bias/volatility leads to longer sampling. Indeed, it holds that $|Z_t^{LH}| > |Z_t^{LL}|$ (0.130 vs. 0.058) and $|Z_t^{HH}| > |Z_t^{HL}|$ (0.11 vs. 0.078), both differences

[5]The trimmed mean is a measure of central tendency which involves the calculation of the mean after discarding parts at the high and low ends of a distribution. We use this measure, along with the standard deviation, to mitigate the impact of outliers.

significant at the 1% level (p<0.000)[6].

The above results can be also confirmed by a simple regression. Table 8 reports the treatment effects. As dependent variable we use the absolute value of the threshold $Z_t$, with explanatory variables the repetition number and treatment dummies for each of the treatments. We set the treatment LL as the default. The coefficients can be interpreted as the effect that the various treatments have on the absolute level of the threshold. All the coefficients, except the repetition coefficient, are significant and have the signs that theory predicts (with the exception of the HL, where one would expect a negative sign).

Finally, before moving to the structural estimation, it would be interesting to investigate whether the subjects exploit the acquired information in a fruitful way, that is, whether spending money on sampling, helped them to make the right decision. Table 9 reports the number of trials in which a subject sampled, bet on the ambiguous urn and also bet on the colour with the positive bias in that repetition. The percentage of successful trials ranges from 63 to 81% with the rate in high bias/volatility treatments being significantly greater to that of the low bias/volatility ones.

## 5.2 Structural estimation

To further explore the internal validity of the model, we opt for structural parametric estimations of the two models. As our dataset is not balanced (some subjects sample for the whole 20 repetitions, while others for very few), and as the number of observations per subject does not allow for model-fitting at the subject-level, we pool all the data together and we estimate mixture specifications (Harrison and Rutström

[6]The results on Hypotheses 4 and 5 are robust when the comparison are based on sampling time, or when the comparison is based on the maximum reached sampling threshold (rather than the threshold level when the subject pushed the stop button).

2008; Bruhin et al. 2010; Conte et al. 2011). In particular, we estimate a structural finite mixture model assuming two data-generating processes, a standard Bayesian learning model, in which we set the value of $\epsilon$ to be equal to zero, and the Epstein and Ji (2019) model.

The basic idea of a mixture model is to assign each subject's choices to one of the two decision making models. As the model presented in section 3 is deterministic, to aid the econometric estimation, we assume that a decision maker chooses her optimal stopping threshold with some noise. Therefore, for a particular repetition $n$ we assume that the actual stopping threshold $Z_{tn}$ is given by $Z_{tn} = \bar{z}_n + u$, with $\bar{z}_n$ being the optimal stopping threshold and $u$ a Fechnerian error such that $u \sim N(0, \xi)$ and $\xi$ the standard deviation to be estimated.

The likelihood, conditional on the Bayesian model being true, depends only on the observed choices and the value of $\xi$. In particular, the density of this type, for an i-th subject can be expressed as

$$f^B(Z_{tn}, \xi) = \prod_{n=1}^{N} \frac{1}{\xi} \phi \left( \frac{\bar{z}_n - Z_{tn}}{\xi} \right)$$

where $\phi(.)$ denotes the density of the standard normal distribution, and $N$ the total number of repetitions.

The conditional likelihood for the Epstein and Ji (2019) model is defined in a similar way, with the only difference being that on top of the observed choices and the value of $\xi$, it also depends on the parameter of ambiguity aversion $\epsilon$.

$$f^{EZ}(Z_{tn}, \xi, \epsilon) = \prod_{n=1}^{N} \frac{1}{\xi} \phi \left( \frac{\bar{z}_n - Z_{tn}(\epsilon)}{\xi} \right)$$

If we let $\pi_{mix}$ denote the probability that the Bayesian learning model is correct and $1 - \pi_{mix}$ the probability that the Epstein and Ji (2019) model is correct, the log-likelihood of the finite mixture model is then given by

$$\ln L(\xi, \epsilon, \pi_{mix}, Z) = \sum_i \ln \Big( \pi_{mix} \times f^B(Z_{tn}, \xi) + (1 - \pi_{mix}) \times f^{EZ}(Z_{tn}, \xi, \epsilon) \Big)$$

As the conditional likelihoods are defined at the subject level, the mixing proportion $\pi_{mix}$ classifies each subject to either one or the other model[7]. In Table 10 we report two sets of estimates. For the first one, we include all the observations for which a subject sampled for some time and then chose to bet on the ambiguous bag, at least once (96 subjects). This includes subjects that may have experimented for few rounds, discovered that learning is not useful, and switched back to the risky bag. In the second set, we include only the observations of those subjects who sampled for at least 15 repetitions, and bet on the ambiguous bag for 75% or the repetitions or more (44 subjects).

We estimate the model using MLE techniques. When all the data are pooled together, 47% of the subjects are classified as Bayesian, while the choices of 53% of the subjects can be explained by Epstein and Ji (2019). When the constrained sample is used, the proportion of Bayesian subjects drops to 40%. A potential explanation could be that the proportion of Bayesian subjects drops, as we remove the observations of those subjects who sampled for very few repetitions. These subjects for example, could have sampled for very little time for few rounds, which impacts the overall mean time spent on sampling. This can also be explained with the parameter of ambiguity aversion, which is higher in the constrained sample.

[7] An alternative way to set up the overall likelihood would be to follow Harrison and Rutström (2008) and assume that each observation is generated by one of the two models. We opt for the classification at the subject level as it has a more straightforward interpretation.

| Treatment | subjects | sessions | $c$ | $\alpha$ | $\sigma$ | $\alpha/\sigma^2$ |
|---|---|---|---|---|---|---|
| LL | 28 | 2 | 0.02 | 0.15 | 0.150 | 6.67 |
| LH | 25 | 2 | 0.02 | 0.30 | 0.212 | 6.67 |
| HL | 27 | 2 | 0.04 | 0.15 | 0.150 | 6.67 |
| HH | 34 | 2 | 0.04 | 0.30 | 0.212 | 6.67 |
| All | 114 | 8 | - | - | - | - |

Table 1: Details of the 4 treatments LL, LH, HL and HH (the first element indicates the level of cost while the second one the level of $\alpha$).

| Treatment | sampled | trials | % |
|---|---|---|---|
| LL | 358 | 560 | 0.639 |
| LH | 264 | 500 | 0.528 |
| HL | 364 | 540 | 0.674 |
| HH | 483 | 680 | 0.710 |
| All | 1469 | 2280 | 0.644 |

Table 2: Percentage of trials in which subjects sampled.

| Treatment | Choose A | trials | % |
|---|---|---|---|
| LL | 105 | 202 | 0.520 |
| LH | 184 | 236 | 0.780 |
| HL | 136 | 176 | 0.773 |
| HH | 176 | 197 | 0.893 |
| All | 601 | 811 | 0.741 |

Table 3: Percentage of trials in which subjects did not sample and chose the risky Bag.

| Treatment | Choose A | trials | % |
|---|---|---|---|
| LL | 73 | 358 | 0.204 |
| LH | 24 | 264 | 0.091 |
| HL | 70 | 364 | 0.192 |
| HH | 42 | 483 | 0.087 |
| All | 209 | 1469 | 0.142 |

Table 4: Percentage of trials in which subjects sampled and chose the risky Bag.

| Treatment | Mean | Mean_10 | Median | Sd |
|---|---|---|---|---|
| LL | 0.058 | 0.050 | 0.043 | 0.049 |
| LH | 0.130 | 0.115 | 0.094 | 0.106 |
| HL | 0.078 | 0.070 | 0.068 | 0.056 |
| HH | 0.110 | 0.099 | 0.089 | 0.083 |
| All | 0.094 | 0.082 | 0.072 | 0.081 |

Table 5: Threshold level: mean, trimmed mean at 10%, median and standard deviation.

| Treatment | Mean | Mean_10 | Median | sd |
|---|---|---|---|---|
| LL | 9.730 | 7.530 | 6.190 | 10.340 |
| LH | 17.790 | 16.230 | 13.790 | 14.500 |
| HL | 14.890 | 13.230 | 11.690 | 11.820 |
| HH | 14.490 | 12.490 | 14.490 | 12.710 |
| All | 14.136 | 12.100 | 9.990 | 12.650 |

Table 6: Sampling time in seconds: mean, trimmed mean at 10%, median and standard deviation.

| Treatment | sampled | sampled all 20 | % |
|---|---|---|---|
| LL | 110 | 140 | 0.785 |
| LH | 153 | 160 | 0.956 |
| HL | 167 | 220 | 0.759 |
| HH | 346 | 380 | 0.910 |
| All | 776 | 900 | 0.862 |

Table 7: Percentage of trials in which subjects sampled and chose the ambiguous bag (includes only subjects who sampled in all 20 rounds).

| | Estimate | Std. Error | t value | Pr(>∣t∣) | |
|---|---|---|---|---|---|
| (Intercept) | 0.054 | 0.006 | 9.062 | 0.000 | *** |
| dHH | 0.052 | 0.006 | 8.939 | 0.000 | *** |
| dHL | 0.020 | 0.006 | 3.181 | 0.002 | ** |
| dLH | 0.072 | 0.007 | 10.686 | 0.000 | *** |
| round | 0.000 | 0.000 | 0.929 | 0.353 | |
| $R^2$=0.103, N=1255 | | | | | |

Table 8: Treatment effects including only the trials in which subjects sampled and chose the ambiguous bag. Dependent variable: absolute value of threshold $Z_t$. dHH is a dummy variable for treatment HH and so on. LL is the baseline. .

| Treatment | Correct colour | trials | % |
|---|---|---|---|
| LL | 179 | 284 | 0.630 |
| LH | 193 | 239 | 0.808 |
| HL | 202 | 294 | 0.687 |
| HH | 324 | 438 | 0.740 |
| All | 898 | 1255 | 0.716 |

Table 9: Percentage of trials in which subjects sampled, chose the ambiguous bag and chose the correct colour.

| Parameter | (1) | (2) |
|---|---|---|
| $\epsilon$ | 0.309 | 0.351 |
| s.e. | 0.007 | 0.009 |
| $\xi$ | 0.070 | 0.080 |
| s.e. | 0.002 | 0.002 |
| $\pi_{mix}$ | 0.473 | 0.404 |
| s.e. | 0.048 | 0.072 |
| $\ln L$ | 1391.13 | 800.03 |
| Parameters | 3 | 3 |
| Observations | 1249 | 776 |

Table 10: Maximum likelihood estimates of the finite mixture model. Model (1) includes all the observations in which the subjects sampled and bet on the ambiguous urn at least once. Model (2) includes all the observations in which the subjects sampled for more than 15 rounds and bet on the ambiguous urn for more than 75% of those repetitions.

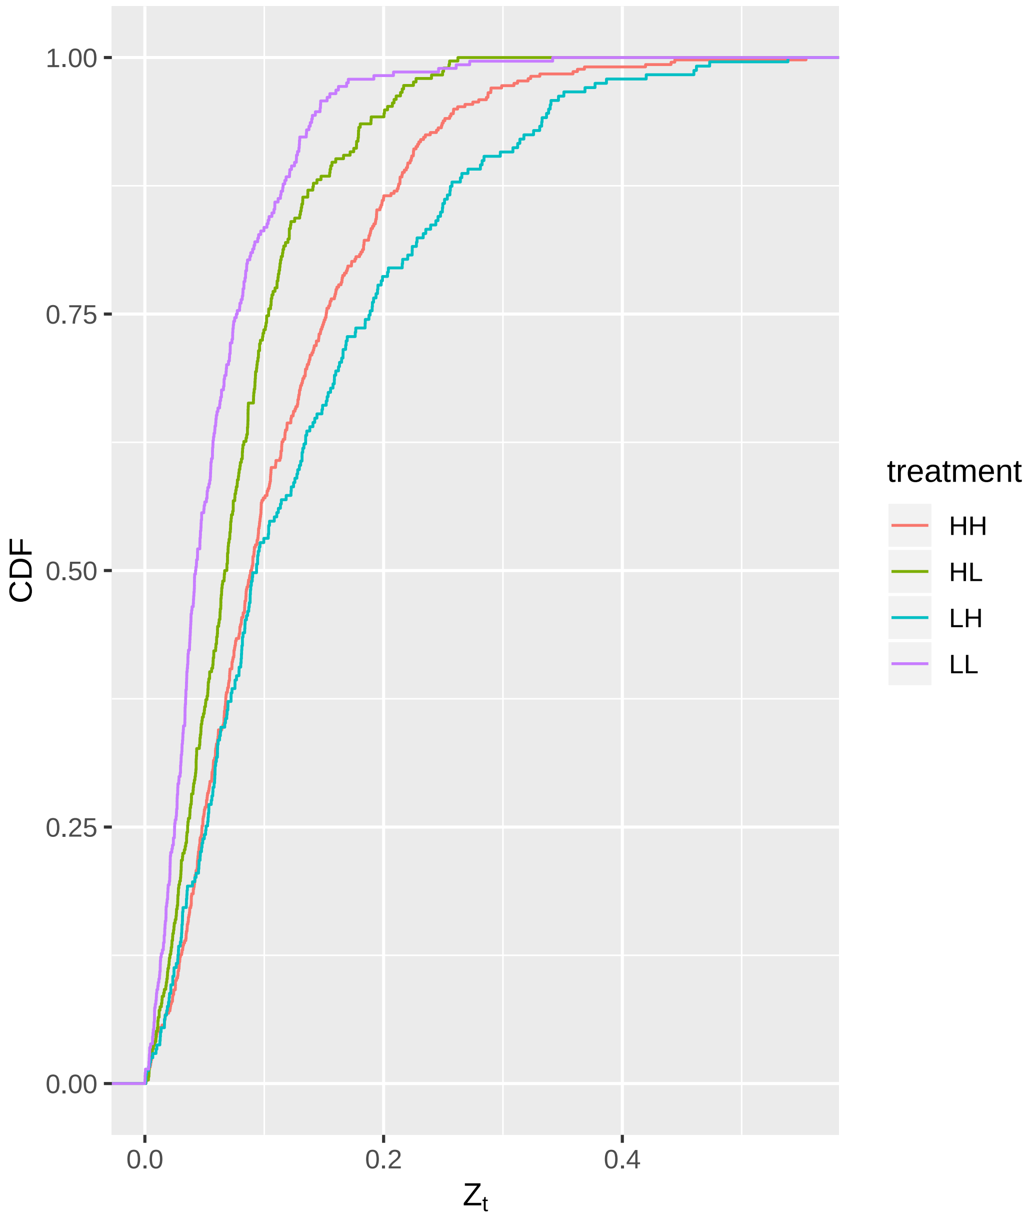


Figure 1: Empirical CDFs of observed stopping thresholds.

# References


Baillon, A., H Bleichrodt, U. Keskin, O. L'Haridon, and C. Li (2018), "Learning Under Ambiguity: An Experiment Using Initial Public Offerings on a Stock Market." *Management Science*, 64, 2181–2198.

Bock, O., I. Baetge, and A. Nicklisch (2014), "hroot: Hamburg Registration and Organization Online Tool." *European Economic Review*, 71, 117 – 120.

Bruhin, A., H. Fehr-Duda, and T. Epper (2010), "Risk and rationality: Uncovering heterogeneity in probability distortion." *Econometrica*, 78, 1375–1412.

Conte, Anna, John D. Hey, and Peter G. Moffatt (2011), "Mixture models of choice under risk." *Journal of Econometrics*, 162, 79 – 88.

Engle-Warnick, J. and S. Laszlo (2017), "Learning-by-doing in an Ambiguous Environment." *Journal of Risk and Uncertainty*, 71–94.

Epstein, L. and S. Ji (2019), "Optimal Learning under Robustness and Time-Consistency." *Operations Research*.

Ert, E. and S. Trautmann (2014), "Sampling Experience Reverses Preferences for Ambiguity." *Journal of Risk and Uncertainty*, 49, 31–42.

Fudenberg, D., P. Strack, and T. Strzalecki (2018), "Speed, Accuracy, and the Optimal Timing of Choices." *American Economic Review*, 108, 3651–84.

Harrison, G. and E. Rutström (2008), "Expected Utility Theory and Prospect Theory: one Wedding and a Decent Funeral." *Experimental Economics*, 12, 133.

Nicholls, Nicky, Aylit Tina Romm, and Alexander Zimper (2015), "The Impact of Statistical Learning on Violations of the Sure-thing Principle." *Journal of Risk and Uncertainty*, 50, 97–115.

Trautmann, S. and R. Zeckhauser (2013), "Shunning Uncertainty: The Neglect of

Learning Opportunities." *Games and Economic Behavior*, 79, 44 – 55.