Learning Under Uncertainty with Multiple Priors: Experimental Investigation*

April 30th, 2021

James Bland† Yaroslav Rosokha‡

Abstract

We run an experiment to compare belief formation and learning under ambiguity and under compound risk at the individual level. We estimate a four-type mixture model assuming that, for each type of uncertainty, subjects may either learn according to Bayes' Rule or learn according to a multiple priors model of learning. Our results indicate that majority of subjects are Bayesian, both under compound risk and under ambiguity, while the second most frequent type are subjects that are Bayesian under compound risk but who use a multiple priors model of learning under ambiguity. In addition, we find strong evidence against a common assumption that participants' initial beliefs (and priors) are consistent with information provided about the uncertain process.

JEL classification: *C91, D83*

Keywords: Experiments, Learning, Ambiguity, Compound Risk, Multiple Priors, Mixture Models

† Department of Economics, The University of Toledo • ORCID iD:https://orcid.org/0000-0002-7117-9998 • Email:james.bland@utoledo.edu

‡ Krannert School of Management, Purdue University • ORCID iD:https://orcid.org/0000-0002-8567-659X • Email:yrosokha@purdue.edu

∗ This paper benefited from discussions and comments from Ernan Haruvy, Othon Moreno, Dale Stahl, as well as from workshop participants at the 2012 North American ESA meetings, the 2014 Annual Informs Meetings, and seminar participants at the University of Texas and Purdue University. This work was supported by a grant from the Russell Sage Foundation.

1 Introduction

Learning from a signal when the initial information is uncertain is critical for economic success. For example, in an innovation context, firms must decide whether to continue with an R&D project depending on research results; in a consumer choice context, people must decide whether to buy a product, based on the product review; in a retail context, managers must decide which assortment of products to offer depending on the product sales. However, in most situations, a decision maker does not fully understand, has little information about, or considers multiple theories about the process generating the signal. In such cases, it is common practice to model the environment as uncertain and use Bayes' rule as a way for the agent to learn from a signal. Little is known, however, regarding the learning process under uncertainty with unknown probabilities (henceforth *ambiguity*). In particular, is it different from the learning process under uncertainty with known probabilities (henceforth *compound risk*)? And, how well does Bayes' rule capture the learning process under ambiguity?

The difference between ambiguity and risk was first noted by Knight (1921). Later, using a thought experiment, Ellsberg (1961) showed that behavior under ambiguity cannot be explained by the subjective expected utility theory of Savage (1954). Recent experimental studies show that there is substantial heterogeneity in attitudes towards compound risk and/or ambiguity at the individual level (Halevy, 2007; Stahl, 2014; Abdellaoui, Klibanoff, and Placido, 2015; Harrison, Martínez-Correa, and Swarthout, 2015). In addition, a number of studies find an association between attitudes towards the compound risk and ambiguity (Halevy, 2007; Abdellaoui, Klibanoff, and Placido, 2015; Dean and Ortoleva, 2015; Prokosheva, 2016; Qiu and Weitzel, 2016; Chew, Miao, and Zhong, 2017). While the above studies focused on decision-making in static environments, there is scarce evidence on any such relationship about learning under compound risk and ambiguity.

In this paper, we present an experiment designed to compare the learning process under compound risk and under ambiguity at the individual level. In our experiment, there are two types of urns composed of black and white marbles. Compound risk urns are constructed by randomly drawing from a set of urns with known composition. Thus, the subjects are provided with the objective prior about the probability that a black (or white) marble could be drawn. Ambiguous urns are constructed so that subjects do not know the exact composition of the urn, but know the total number of marbles, which is kept the same as in the compound case. In other words, subjects are not provided with enough information to form an objective prior. In the experiment, each subject faces decisions regarding both types of urn. The questions that we address in this paper deal with the priors considered by the subjects and the process by which those priors are updated. In particular, the following questions are of interest: Are beliefs consistent with the urn composition process? Are there behavioral differences between learning under the compound risk and ambiguity? Can a multiple priors approach explain the learning behavior under ambiguity and/or compound risk?

To answer these questions, we use a mixture model to estimate the proportion of subjects that learn according to Bayes' rule and the proportion of subjects that learn according to a more general, multiple-priors model of Epstein and Schneider (2007). The multiple-priors model fits within a stream of literature that uses maximum likelihood as a way to discriminate among priors after a signal has been observed (Gilboa and Schmeidler, 1993). In particular, agents consider multiple priors about the signal generating process, and upon realization of the signal agents evaluate which of the priors were "likely" to generate the signal. Then only these "likely" priors are updated according to Bayes' rule and considered as decision relevant. Importantly, the model can be applied to the compound risk environment – in which case only the single objective prior will be in the set. We find that the majority (60%) of subjects are Bayesian both under compound risk and ambiguity. We also find a substantial fraction (25%) of subjects who are Bayesian under compound risk but not under ambiguity.

The challenge in considering the multiple priors model is that the set of possible priors is infinite. We estimate two models with different assumptions on the type of priors subjects may use. The first model assumes that subjects' priors are over the possible urns that could be generating the signals (i.e. priors over the number of black vs. white marbles in the urn). We refer to these priors as *Simplex* priors, because they take the form of an element of the 3-dimensional simplex. The second model assumes that subjects' priors take on a *Beta* distribution over the probability that a black marble is drawn. This second class of priors was recently used by Moreno and Rosokha (2016) as part of a behavioral model of belief updating. We find that under the assumption of Simplex priors participants' behavior is in line with the multiple-priors model under both compound risk and ambiguity. At the same time, under the assumption of Beta priors the behavior is more in line with subjects being Bayesian. Our model selection result, however, provides overwhelming evidence in favor of the Beta priors. In particular, while the Simplex priors correspond to the possible urn compositions (and are consistent with the information provided about the uncertain process), they impose implicit restrictions on the strength of priors and the range of the beliefs that subjects could hold. These restrictions prove to be too limiting in describing human belief formation and learning processes as compared to a set of more general Beta priors.

Our work contributes to the literature that investigates learning under compound risk and/or ambiguity. In particular, there exists a large body of literature in economics and psychology with focus on learning under compound risk. The conclusions in this literature vary. For example, in a seminal article, Kahneman and Tversky (1973) present evidence that individuals over-value new information relative to Bayes' rule (a judgment bias known as representativeness). At the same time, other studies (e.g., Buser, Gerhards, and Van Der Weele, 2018; Coutts, 2019) find that subjects under-value new information relative to Bayes' rule (a judgment bias known as conservatism), or that most behavior is well described by Bayes' rule (e.g., El-Gamal and Grether, 1995).

A smaller stream of literature has focused on learning under ambiguity (Cohen, Gilboa, Jaffray, and Schmeidler, 2000; Dominiak, Dürsch, and Lefort, 2012; Baillon, Bleichrodt, Keskin, l'Haridon, and Li, 2013; Qiu and Weitzel, 2013; Ert and Trautmann, 2014; Moreno and Rosokha, 2016). In this literature, the most closely related study to the current paper is Moreno and Rosokha (2016) who develop a behavioral model of belief updating and then estimate their model at the aggregate level. The authors find that learning under compound risk is consistent with Bayes' rule, while the learning process under ambiguity is consistent with over-weighting of the new signal. In the current paper we differ in several important ways: First, we use a within-subject design which allows us to address learning by the same individual in the two environments. Second, we consider a multiple-priors model of learning developed for ambiguous environments, rather than using a reinforcement type behavioral model. Third, we investigate two different specifications of subjective priors. Finally, we estimate a mixture model of different types allowing for individual level heterogeneity in preference, learning, and precision parameters.

The rest of the paper is organized as follows. In Section 2, we describe the experimental design and elicitation procedure and present an overview of the data. In Section 3, we present the learning model and estimation procedure used. In Section 4, we present and discuss our main results. Finally, in Section 5, we conclude.

2 Experimental Design

Design of the current experiment builds on the work by Moreno and Rosokha (2016) to allow for estimation of the multiple priors model of learning and to allow for comparison of learning between compound risk and ambiguity at the individual level. In particular, similar to the prior work, compound risk and ambiguity are implemented using urns of black and white marbles (Figure 1).

Figure 1: Urn Types

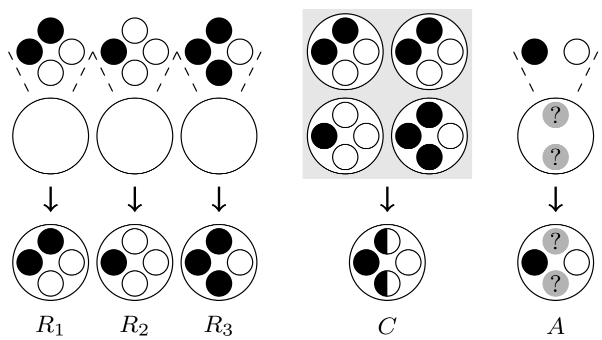

Notes: R_i - risky urn. C - compound urn. A - ambiguous urn. Urns $R_1 - R_3$ are constructed in front of the participants. Urn C is determined by randomly drawing one of the four urns that were constructed in front of the participants. Urn A is constructed by placing two marbles in the urn before subjects enter the room. During the experiment, subjects verify that there are two marbles in the urn, and they are informed that each could be either black or white, but they are not informed about the process by which the marbles were selected; then, one black and one white marble are added to the urn in front of the participants.

The three types of urns used in the experiment differ with respect to their composition process. Specifically, subjects see the exact composition of the risky urns (R_1, R_2, and R_3 in Figure 1). Subjects see the composition "process" of the compound urn (C in Figure 1). Subjects do not know the composition process of the ambiguous urns (A in Figure 1). Thus, no objective probabilities are provided during composition of the A urn, and, therefore, there is ambiguity about the number of black and white marbles. Nevertheless, the same three compositions are possible under the C and A scenarios.

Decision tasks in the experiment involve choosing between a lottery involving one of the urns presented in Figure 1 and a sure option. Specifically, Figure 2 presents the Multiple Price List design (Holt and Laury, 2002; Harrison and Elisabet Rutström, 2008) that we implemented in each task of the experiment.[1] Note that unlike the work by Moreno and Rosokha (2016), tasks in the current experiment involve decisions about black and white marbles. That is, we used two sets of decision tasks: one containing questions of the form "Please choose between \$X (Option A) vs. \$33 if black is drawn, \$5 otherwise (Option B)" and the other containing questions of the form "Please choose between \$X (Option A) vs. \$33 if white is drawn, \$5 otherwise (Option B)." The two types of decision tasks allow for estimation of a set of priors.

Figure 2: Decision Tasks

The outcome of the lottery is based on the color of the ball that will be drawn from urn i. Please choose between Options A and B for each question.

	Option A	**Option B**
1)	\$ 6	\$33 if black (white), \$5 otherwise
2)	\$ 9	\$33 if black (white), \$5 otherwise
3)	\$ 11	\$33 if black (white), \$5 otherwise
4)	\$ 14	\$33 if black (white), \$5 otherwise
5)	\$ 17	\$33 if black (white), \$5 otherwise
6)	\$ 20	\$33 if black (white), \$5 otherwise
7)	\$ 23	\$33 if black (white), \$5 otherwise
8)	\$ 26	\$33 if black (white), \$5 otherwise
9)	\$ 30	\$33 if black (white), \$5 otherwise

Notes: In the experiment, we used two sets of decision tasks. One set had questions of the form "\$X vs. \$33 if black is drawn, \$5 otherwise" and the second set of the form "\$X vs. \$33 if white is drawn, \$5 otherwise." For each decision task, participants had to make the choice for each of the nine questions (with multiple switching points allowed).

[1] While the gap between the certain amounts for Option A (Figure 2) may seem large, it is worth noting that in the experiment subjects will see two signals about each urn with each signal consisting of three draws with replacement. For context, the indifference point for a risk neutral subject who considers objective composition of the C urn, who follows Bayes' rule, and who sees six successes would increase by approximately \$6.

The goal of the current experiment is to allow comparison of the learning process between compound risk and ambiguity at the individual level; therefore each subject is presented with both the decision task involving the compound urn (henceforth C-task) and the decision task involving the ambiguous urn (henceforth A-task). In order to ensure that the order of presentation does not affect the learning process we ran the experiment using two order treatments as presented in Table C-2 of the Online Appendix.

Figure 3 presents the summary of the draws and treatments in the experiment. In total, we recruited eighty-four undergraduate students for the experiment at the University of Texas at Austin.[2] Ten sessions of the experiment were administered between October of 2012 and February of 2013. Each participant made either 144 or 180 decisions over a period of approximately 45 minutes. At the end of the experiment, two decisions were picked at random and carried out to determine the participants' earnings for the experiment.[3] All lotteries were executed by physical randomization devices.

[2]For robustness, we augment the current data set with data from Moreno & Rosokha (2016) that includes 113 participants. Results are presented in Online Appendix B.

[3]The random lottery incentive mechanism has several known issues when it comes to the elicitation of preferences for risk and uncertainty. For example, Freeman, Halevy, and Kneeland (2019) show that when compensated based on one randomly selected lottery from a list, subjects are more likely to select the sure payment over risky lottery as compared to the case when facing only one payoff-relevant decision. Harrison, Martínez-Correa, and Swarthout (2015) show that the random lottery incentive mechanism may in itself lead to violations of the reduction of compound lotteries. However, a between-subject design with each subject facing only one, payoff-relevant decision would not be able to address whether the same subject learns differently under compound risk or under ambiguity.

Figure 3: Experiment Summary

Session	N	Stage 2	Signal 1	Signal 2	Stage 3	Signal 1	Signal 2	Stage 4	Signal 1	Signal 2	Av.Earn
1	11	C	●●●	●●●	A	○○○	○○○	–	–	–	48.00
2	7	A	○○○	●○●	C	●○●	○●●	–	–	–	38.15
3	4	A	●●○	○●●	C	●●●	●●○	A	●●●	●●●	45.00
4	11	A	●○●	●●○	C	○○●	●●●	A	○●○	○○●	40.00
5	11	C	○○●	●○●	A	○●○	○●○	A	○●○	●○○	38.40
6	6	C	○○●	●○●	A	○○○	●●●	A	○○●	●●○	41.83
7	10	A	●●●	●●●	C	●○○	●●○	A	○○○	●●○	51.88
8	10	A	●○○	●●●	C	○○○	●●●	A	○●●	○○●	46.00
9	6	C	●●●	●●●	A	○○○	●○●	A	○●○	○●●	30.17
10	8	C	○○○	●○●	A	○○○	○○●	A	○○○	●○○	42.88
Total	84										42.42

Notes: Stage 1 (not shown) involved tasks regarding the risky (R) urns. No draws were made from the R urns. Stage $i \in \{2, 3, 4\}$ contained a sequence of tasks, each involving either the compound (C) or the ambiguous (A) urn. The task and draws were arranged as follows: first, subjects faced two tasks – one regarding black and one regarding white – before any draws were made; next, three draws with replacement were made (Signal 1), after which subjects again faced the two tasks presented in Figure 2; finally, three more draws with replacement were made (Signal 2), after which subjects faced two more tasks.

3 Learning With Multiple Priors

We consider the MP model developed by Epstein and Schneider (2007). In this model, the parameter of interest is $0 \leq \alpha \leq 1$, which determines the extent to which the decision maker discards "unlikely" priors. In particular, only priors considered as "likely" are updated according to Bayes' rule after each draw. The extent to which priors are "likely" is determined by the likelihood-ratio test relative to the prior with the highest likelihood. Among the obtained posteriors, the one that yields the worst expected payoff is considered for lottery evaluation. Specifically, let M_0 be the set containing all the considered priors at round zero. After observing a history of draws H_{t-1}, the likelihood of each prior $\mu_0 \in M_0$ is evaluated. Then, the decision maker discards all priors μ_0 that do not pass a likelihood-ratio test against an alternative theory that puts maximum likelihood on the sample. Posteriors $\mu_t(H_{t-1}; \mu_0)$ are formed only for priors that pass the test. Thus, the set of posteriors is given by

$$M_t = \{\mu_t(H_{t-1}; \mu_0) : \mu_0 \in M_0, L(H_{t-1}|\mu_0) \geq \alpha \times \max_{\mu'_0} L(H_{t-1}|\mu'_0)\}. \quad (1)$$

Equivalently, a posterior $\mu_t(H_{t-1}; \mu_0)$ will be included in the set of posteriors if $\frac{L(H_{t-1}|\mu_0)}{L(H_{t-1}|.)} \geq \alpha$, where $L(H_{t-1}|.)$ is the highest likelihood observed for all priors in the original set. With the new set M_t, agents make their decision according to the *maxmin* criterion, as in Gilboa and Schmeidler (1989). Notice that M_{t+1} is constructed from M_0, the set of considered priors at round zero, and not from M_t. Note, also, that the higher the α, the smaller the set of posteriors at every t.

The difficulty with estimating the priors is that the types of priors can vary greatly. In order to facilitate the estimation of the sets of priors, we limit our attention to two classes of priors that subjects may use. The first class, which we term *Simplex* priors, is characterized by the probabilities assigned to each of the three compositions of the C and the A urns which are possible in the experiment (and participants know that). The second class, which we term *Beta* priors, is characterized by the belief about the probability of a black marble occurring and the strength of that belief. Next, we describe the two types of priors in more detail.

3.1 Simplex Priors

The first class of priors is motivated by the urn composition process. Specifically, in both the compound and ambiguous scenarios, there are three possible states of the word each corresponding to one of the three urns: R_1, R_2, and R_3. That is, the first class of priors is obtained by assuming that subjects' priors are over the three possible urns that could be generating the signals. We call this class the Simplex priors, as they take the form of an element of the 3-dimensional simplex. Any prior of this form can be parameterized by the three probabilities assigned to each urn, $\mu_0 \in \Delta^3$. So in order to estimate the multiple-priors model with Simplex priors, we need to estimate the set of priors $M_0 = \{\mu_0 : \mu_0 \in \Delta^3\}$. Following history H_{t-1}, a Bayesian using a simplex prior μ_{t-1} updates their beliefs after observing a_t white marbles and b_t black marbles as follows:

$$\mu_{t,k} \propto \binom{a_t + b_t}{b_t} \left(\frac{k}{4}\right)^{b_t} \left(\frac{4-k}{4}\right)^{a_t} \mu_{t-1,k}, \quad k = 1, 2, 3 \tag{2}$$

where $\mu_{t,k}$ is the subject's belief that there are k black marbles in the urn. Noting that the binomial coefficient in (2) is not a function of k, beliefs can be normalized to sum to one as follows:

$$\mu_{t,k} = \frac{k^{b_t}(4-k)^{a_t}\mu_{t-1,k}}{\sum_{l=1}^{3} l^{b_t}(4-l)^{a_t}\mu_{t-1,l}}, \quad k = 1, 2, 3 \tag{3}$$

When making decisions in the A- and C-tasks, the information needed from these beliefs is the posterior probability of drawing a black marble. This is equal to:

$$p_t = \sum_{k=1}^{3} \frac{k}{4}\mu_{t,k} \tag{4}$$

3.2 Beta Priors

The second class of priors is motivated by a behavioral model developed in Moreno and Rosokha (2016). In this class, subjects' priors take on a Beta distribution over the probability that a black marble is drawn. Specifically, prior, $P(p|H_t)$, is distributed according to a Beta distribution with parameters a_t, and b_t. The properties of the Beta distribution imply that the history, H_t, is summarized by (a_t, b_t). Furthermore, after observing a signal, s_t, the posterior is distributed according to Beta distribution with parameters $a_{t+1} = a_t + s_t$, $b_{t+1} = b_t + (3 - s_t)$. A transformation that facilitates interpretation of Beta priors is that the Beta distribution can be equivalently characterized by the mean $p_t = \frac{a_t}{a_t+b_t}$ and the strength $N_t = a_t + b_t$. And so, in order to estimate the MP model with Beta priors, we need to estimate the set of priors $M_0 = \{\mu_0 : \mu_0 \sim Beta(p_0, N_0)\}$ and the power of the likelihood-ratio test, α.

3.3 Examples

In order to better understand the model and the two different assumptions regarding the priors, we consider the following examples for a hypothetical sequence of draws from the A urn. Specifically, suppose that before any draws have been made, an agent considers the set of priors, M_0, given by $\delta \in \left(0, \frac{1}{3}\right)$, where δ is the minimum probability that the subject assigns to each possible urn composition. Specifically, we assume that the set of priors can be characterized by the subset of the simplex that places at least probability mass δ on each urn:

$$M_0 = \{\mu_0 : \mu_0 \in \Delta^3 \text{ and } \min\{\mu_0\} \geq \delta\}. \tag{5}$$

Figure 4 presents the mechanics of the learning process. Specifically, a subject starts out with a set of priors that he considers before any draws have been made, M_0, which contains all priors that place at least $\delta = .1$ on each possible urn composition (Figure 4 (a)). By eliciting his beliefs about the probability of a black and a white marble being drawn we are able to pin down the best- and the worst-case scenarios.

Figure 4: Learning with Simplex Priors

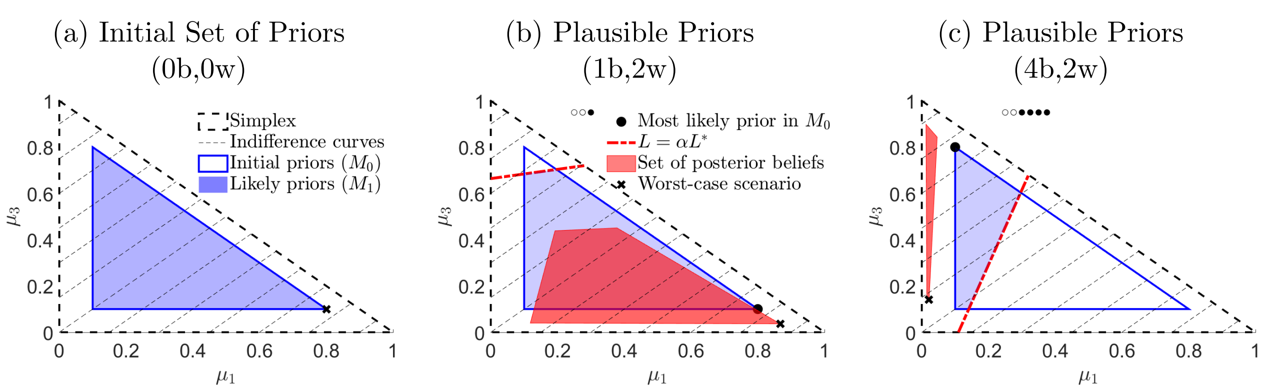

Notes: An example of learning with simplex priors for $\alpha = 0.4$ and $\delta_1 = \delta_2 = \delta_3 = 0.1$, where α is the parameter of the multiple-priors model of Epstein and Schneider (2007) and δ_i is the minimum probability that the subject assigns to each possible urn composition.

Suppose that one black and two white marbles are drawn from the urn, then learning can be summarized in three steps: First, the agent determines the likelihood of each prior generating the sequence. Second, the agent keeps only the priors (Figure 4 (b)) that pass the likelihood ratio test (the set of priors bound by the solid red line in Figure 4 (b)). Third, the agent forms a posterior for each of the "likely" priors (Figure 4 (b) red shaded area). Then, the worst-case scenarios, with respect to the probability of a black marble being drawn, are the posterior beliefs that minimize the probability of drawing a black marble. Inspection of Figure 4 (b) reveals that, following our assumption about the shape of M_0, the worst-case scenario must lie at one of the corners of the set of posterior beliefs. This is because the agent's indifference curves are linear. Hence, when we take this model to the data, we only need to compute the posterior probability of drawing a black marble at the corners of the red shaded region: if there are multiple posterior beliefs that are equally bad for the agent, they must all imply the same probability of drawing a black marble. Figure 4 (c) shows the sets of likely priors and posterior beliefs after drawing four black and two white marbles. In this scenario, the agent discards priors that assign too much probability mass to the urn with one black marble. The set of posterior beliefs now assign very little probability to the urn having one black marble.

Figure 5 (a) presents the evolution of beliefs (means of priors) corresponding to the priors in Figure 4. In addition, Figure 5 (b) presents an example for the case of Beta priors. Specifically, suppose that before any draws have been made, an agent considers the set of priors, M_0, given by

$$M_0 = \{\mu_0 : \mu_0 \sim Beta(p_0, N_0), p_0 \in [.325, .675], N_0 = 10\}. \tag{6}$$

That is, this agent considers a worst-case scenario about the probability of black and white marbles that are drawn being the same and equal to $p_0(B) = p_0(W) = .325$,[4] and the strengths of all priors in this set being the same and equal to 10.00.

[4]This corresponds to the same ranges of $p_0(B)$ and $p_0(W)$ in the Simplex priors example.

Figure 5: Belief Updating with Simplex and Beta priors

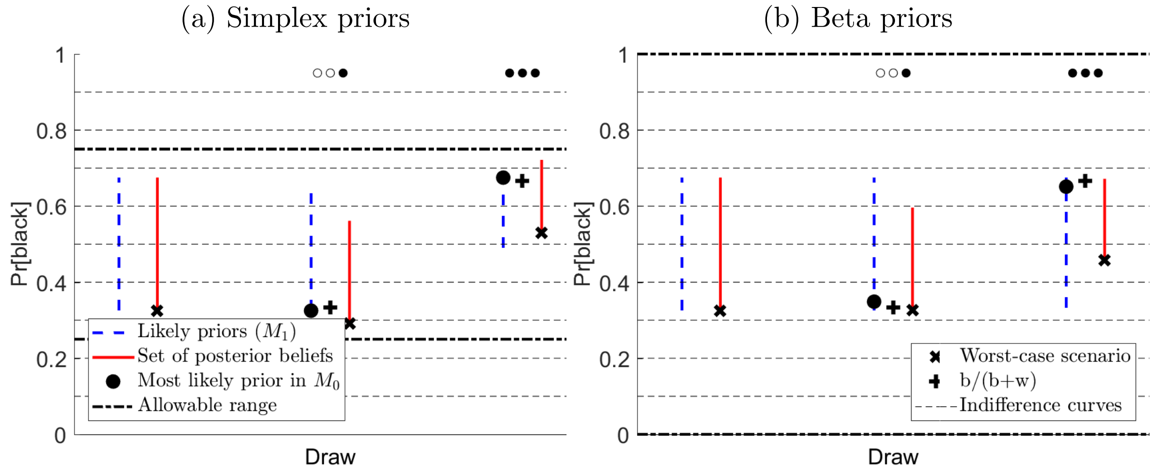

Notes: $\alpha = 0.4$, $\underline{p}_0 = 0.325$, $\overline{p}_0 = 0.675$, $N_0 = 10$. Blue lines show the (remaining) set of prior probabilities of drawing a black marble, red lines show the posteriors of these. Black dot shows the most likely prior, + symbol shows the fraction of black marbles drawn. Black lines show the range of probabilities subject's beliefs must fall into. That is, suppose we observed many, many black signals, then all posteriors would converge from below on $\Pr[black] = 0.75$, even as $\frac{b}{b+w} \to 1$. For Beta priors, posterior $\Pr[black] \to \frac{b}{b+w}$. Therefore the Beta model allows subjects to have beliefs outside of $p \in [0.25, 0.75]$, while Simplex does not.

Note that for a fixed N_0, as α increases, the size of the consideration set, M_t, decreases, leading to a smaller difference between the worst- and best-case scenarios at each point in time. And, as N_0 increases, the effect is the opposite — the worst- and best-case scenarios get further apart.

3.4 Estimation

The novel feature of this paper is that we estimate a mixture model of different types allowing for individual level heterogeneity in preference, learning, and precision parameters. Specifically, we use a hierarchical Bayesian approach to estimate the fraction of subjects that are of Bayesian (B) type (i.e., those who follow Bayes' rule with a subjective prior even when an objective one exists); and the fraction of subjects that are of Multiple priors (M) type (i.e., those who behave according to the multiple priors model of Epstein and Schneider (2007)). We allow for subjects to be of different types in each of the two tasks, so we estimate a four-type mixture model, where each subject is exactly one of (CB, AB), (CB, AM), (CM, AB), or (CM, AM). Next, we describe the estimation procedure in more detail.

To begin, we assume that individuals' utility function is parameterized using a normalized version of the CRRA utility representation:

$$u_i(x) = \frac{x^{1-\gamma_i} - 1}{1 - \gamma_i}, \tag{7}$$

where x is the outcome and γ_i is the risk-aversion parameter to be estimated. Thus, $\gamma_i = 0$

corresponds to subject i being risk-neutral, and $\gamma_i > (<)0$ corresponds to a risk-averse (risk-loving) subject. We use the contextual utility approach of Wilcox (2011) and assume that the agents perceive that the difference between choices is relative to the range of outcomes found in the pair of options. That is,

$$U_i(A) - U_i(B) = \frac{E[u_i(A)] - E[u_i(B)]}{u_i(\$33) - u_i(\$5)}. \tag{8}$$

Notice that \$33 is the best possible outcome and \$5 is the worst possible outcome for all decisions in our experiment. Subject i chooses the option with the highest expected value given her current belief, subject to an error, which is assumed to be distributed according to a logistic distribution centered at zero:

$$P_{A_{i,t}} = \frac{1}{1 + e^{-\lambda_i E_i[U_i(A_{i,t}) - U_i(B_{i,t})]}}, \tag{9}$$

where $P_{A_{i,t}}$ is the probability that the subject chooses option A at round t for the ith lottery pair; $A_{i,t}$ and $B_{i,t}$ are the ith lottery pair presented to the participants in round t; the subscript on the expectation, $E_i[\cdot]$, indicates that subject i is evaluating her utility based on individual-specific parameters, which describe her prior if she is Bayesian, and her α_i and parameters governing her set of priors if she behaves according to Epstein and Schneider (2010). $\lambda_i \geq 0$ is Subject i's logistic choice precision: if $\lambda_i = 0$ she will randomize uniformly over her choice set, and the probability of her choosing the option that maximizes her utility is increasing in λ_i.

Combining equations (7), (8), and (9), we formulate the likelihood function of i's choices y_i, conditional on her parameters θ_i:

$$\mathbf{L}_i(\theta_i, \tau_i) = \prod_{i,t} P_{A_{i,t}}^{y_{i,t}} \times (1 - P_{A_{i,t}})^{(1-y_{i,t})}, \tag{10}$$

where θ_i is a vector of all individual level parameters to be estimated and τ_i is a latent categorical variable identifying the model — either Bayesian or Epstein and Schneider (2007)— that she uses to make decisions in each of the decision tasks. We assume that subjects maximize expected utility in the R-task, but could behave according to either of these models in the C-task or the A-task. Hence, there are four possible types that subjects could be classified into, that is: $\{\text{Bayesian}, \text{Multiple priors}\} \times \{\text{A-task}, \text{C-task}\}$. We use $\rho \in \Delta^4$ to denote the categorical distribution over these four types of subject. While we allow structural parameters θ_i to vary by subject, we assume that each subjects' θ_i is an iid draw from a multivariate normal distribution:

$$\theta_i \sim iidN(\beta, \Sigma) \tag{11}$$

We estimate these hyperparameters β, Σ, and ρ jointly with the θ_i's. We combine (10) with priors over parameters β, Σ, and ρ, then simulate the posterior distribution of all parameters described above using techniques outlined in Online Appendix D.[5]

To summarize, the current experiment and estimation extend Moreno and Rosokha (2016) in four important dimensions. First, we implement a within-subject design, where each subject is faced with a set of tasks about the compound urn (C-tasks) and a set of tasks about the ambiguous urn (A-tasks). This will allow us to identify whether a subject behaves differently under compound risk and under ambiguity. Second, beliefs are elicited both about the proportion of black marbles and the proportion of white marbles. This allows us to estimate a multiple-priors model of learning, rather than focus on learning models with singleton priors. Third, we estimate the model under two different assumptions regarding the priors that subjects use. In particular, we compare the model based on priors that are consistent with the urn composition process, with a model based on subjective priors. Fourth, we estimate a mixture model of different types allowing for individual level heterogeneity. That is, we allow preference and learning parameters to vary across subjects.

4 Results

This section is organized as follows. In Section 4.1, we discuss the results for the two assumptions on the types of priors that subjects may use when learning under compound risk and ambiguity. In Section 4.2, we discuss estimation results for the risk-aversion and the sets of priors that subjects consider.

4.1 Model comparison

Recall, that we consider two different assumptions regarding the class of priors subjects consider. The first class is the Simplex priors over the three possible urns that could be generating the signals (i.e. priors over the number of black vs. white marbles in the urn). This class is consistent with the information provided regarding the urn composition. The second class is the Beta priors over the probability that a black marble is drawn. This class is more general in that it allows subjects to consider priors that are not possible given the composition process of the urns.

Table 1 summarizes important features of the posterior distribution of mixing probabilities. Panel (a) of the figure presents the results for the Simplex priors case. When restricting the behavior to Simplex priors, we estimate that approximately 58% of subjects behave according to the multiple-priors model and approximately 22% of subjects are Bayesian both in the A-task and in the C- tasks. Panel (b) of the figure presents the results for the Beta priors model. When considering a set of more general priors, we estimate that approximately 60% of subjects behave

[5] Our specification is therefore a mixture model at the subject level with a “correlated random coefficients” assumption about individual specific parameters. It is therefore more akin to Conte, Hey, and Moffatt (2011) than Harrison and Rutström (2009) in two important ways: firstly, the mixing is at the subject level rather than the decision level, and secondly, parameters θ_i are assumed to be random draws, rather than deterministic functions of observable characteristics. Our estimation differs from Conte, Hey, and Moffatt (2011) in only three notable ways: (i) we consider different behavioral models, hence our likelihood functions are different; (ii) where Conte, Hey, and Moffatt (2011) has two behavioral types, we have four; and (iii) we use Bayesian techniques instead of (simulated) maximum likelihood.

according to the Bayesian model in both the A- and C- tasks, and approximately 25% of subjects behave according to the multiple-priors model in the A-task and Bayesian in the C-task.

Table 1: Summary of mixing probability estimates

(a) Simplex Priors

Mixing probabilities – joint		
AB CB	0.220	(0.077)
AB CM	0.113	(0.061)
AM CB	0.088	(0.069)
AM CM	0.578	(0.090)
Mixing probabilities – marginal		
CM	0.692	(0.098)
AM	0.666	(0.073)
Pr(CM > AM)	0.636	
Prob modal type		
AB CB	0.009	
AB CM	0.000	
AM CB	0.002	
AM CM	0.989	

Marginal log-likelihood: -4539 (mean)

(b) Beta Priors

Mixing probabilities – joint		
AB CB	0.603	(0.061)
AB CM	0.062	(0.035)
AM CB	0.251	(0.051)
AM CM	0.085	(0.032)
Mixing probabilities – marginal		
CM	0.147	(0.047)
AM	0.336	(0.054)
Pr(CM > AM)	0.002	
Prob modal type		
AB CB	1.000	
AB CM	0.000	
AM CB	0.001	
AM CM	0.000	

Marginal log-likelihood: -4337 (mean)

Notes: Panel (a) shows estimates for Simplex priors. Panel (b) shows estimates for Beta priors. Reported mixing probabilities are posterior means (standard deviations). "Prob modal type" reports the posterior probability that each type is the most prevalent in the population. AB-CB, AB-CM, AM-CB, and AM-CM are the four possible types that subjects could be classified into from the set {Bayesian, Multiple priors} $\times$ {A-task, C-task}.

We compare and select one of these estimated models using a Bayes Factor. In our case we calculate a Bayes Factor of approximately 1.7×10^{68} in favor of the Beta priors specification over the Simplex priors specification.

Result 1 *Subjects do not use (sets of) priors consistent with the urn composition process.*

Result 1 highlights the importance of allowing subjective priors and sets of priors that might be inconsistent with the urn construction process. In particular, if subjects know (have verified) that there are three possible urn compositions, then they should hold a Simplex prior (and posterior), which in addition to the restriction that $p \in [0.25, 0.75]$ also places an implicit restriction on the "weight" of the prior. For example, the maximum amount by which the mean of the Simplex prior can change after one draw is 13.4%. In practice, however, subjects may attribute little weight to the initial prior and form posteriors close to the empirical frequency.[6] Beta priors do not impose such

[6] For robustness, we carry out the analyses from above while restricting the priors (and posteriors) to follow Beta distribution truncated to [.25,.75]. Results are presented in Tables B-1, and B-3 of Online Appendix B. We find the main results are quantitatively similar —the proportion of AB-CB type is approximately 60%, and the proportion of AM-CB type is approximately 29%— though individual preference and learning parameters change.

implicit restrictions. Thus, in terms of the probability of drawing a black marble, the Beta priors assumption is a more flexible model, and this added flexibility improves that model's performance.

Given overwhelming support for the Beta priors, our further analysis focuses only on the Beta priors model. In particular, the second panel of Table 1 reports the marginal mixing probabilities for each task. We find that 15% and 34% of subjects use MP decision rules in the C-task and in the A-task, respectively. These numbers are statistically different: the posterior probability that subjects are more likely to be MP in the C-task than in the A-task, is approximately 0.002. Table 2 explores the ordering of mixing probabilities for the Beta priors model in more detail.

Table 2: Ordering of Mixing Probabilities

Type ranking	Posterior probability
AB CB $\geq$ AM CB $\geq$ AM CM $\geq$ AB CM	0.6955
AB CB $\geq$ AM CB $\geq$ AB CM $\geq$ AM CM	0.2959
AB CB $\geq$ AM CM $\geq$ AM CB $\geq$ AB CM	0.0062
AB CB $\geq$ AB CM $\geq$ AM CB $\geq$ AM CM	0.0018
AM CB $\geq$ AB CB $\geq$ AM CM $\geq$ AB CM	0.0003
AM CB $\geq$ AB CB $\geq$ AB CM $\geq$ AM CM	0.0002
AB CB $\geq$ AM CM $\geq$ AB CM $\geq$ AM CB	0.0001

Notes: AB-CB, AB-CM, AM-CB, and AM-CM are the four possible types that subjects could be classified into from the set {Bayesian, Multiple priors}$\times${A-task, C-task}.

Table 2 presents seven most likely (by posterior probability) orderings of mixing probabilities. Since the prior distribution places equal weight on all mixing probabilities (and hence orderings of these), these numbers are all equal to $1/4! \approx 0.04$ in the prior distribution. In our experiment, 99.1% of the posterior probability is placed on the first two rows of this table. These rows agree on the ranking of the two most prevalent types, (AB-CB) and (AM-CB), which account for about 85% of subjects, but disagree on the ordering of the two least likely types.

Result 2 *The two most common types of subjects are: i) Bayesian under both compound risk and ambiguity, and ii) Bayesian under compound risk, but using a multiple-priors learning model under ambiguity.*

Result 2 states that the most common type of subjects are those that are Bayesian in both the ambiguous and compound environments. The second most common type are subjects that are Bayesian in compound environments but behave according to the multiple priors model in the ambiguous task. Next, we present the individual parameter estimates.

4.2 Parameter Estimates

Table 3 presents summary statistics for the individual-level parameter estimates. In particular, the table presents the means and medians for the preference parameter of risk aversion, the precision parameter, and the learning model parameters for the two common types.

Table 3: Parameter Estimates

	COMMON		COMPOUND - BAYES		AMBIGUOUS - BAYES		AMBIGUOUS - MULTIPLE PRIORS			
	γ	$\log \lambda$	$\log N_0$	p_0	$\log N_0$	p_0	α	$\log N_0$	p_{0mid}	p_{0sp}
MEAN										
	0.62	2.71	1.17	0.49	0.91	0.50	0.83	4.80	0.51	0.66
	(0.09)*	(0.13)[a]	(0.21)*	(0.01)[a]	(0.31)*	(0.01)[a]	(0.12)[a]	(0.47)*	(0.03)[a]	(0.11)[a]
VARIANCE & CORRELATION										
γ	2.19[a]									
$\log \lambda$	0.23	8.30[a]								
$\log N_0$	-0.17	0.17	0.67[a]							
p_0	0.01	0.01	-0.01	0.07[a]						
$\log N_0$	0.11	0.83	0.24	0.01	6.10[a]					
p_0	-0.02	-0.05	0.05	0.01	-0.08	0.19[a]				
α	-0.27	-0.63*	0.06	0.01	-0.51	0.06	0.29[a]			
$\log N_0$	-0.24	-0.00	0.13	-0.02	-0.02	0.03	0.17	0.97[a]		
p_{0mid}	-0.04	0.02	0.18	-0.01	0.04	0.02	0.03	0.02	0.06[a]	
p_{0sp}	-0.21	-0.12	0.20	0.02	-0.06	0.01	0.11	0.06	-0.04	0.15[a]

Notes: Table shows posterior means (standard deviations). * indicates that a 95 percent Bayesian credible region does not include zero. [a] indicates that stars are suppressed because these parameters can only be positive. γ - risk-aversion. λ - logit choice precision. N_0 - prior strength. p_0 prior mean. α - likelihood for discarding priors. p_{0mid} midpoint of the set of prior means. p_{0sp} spread of the set of prior means.

There are several results from Table 3 that are worth noting. First, we find a relatively large spread in the set of priors (0.66) associated with the multiple-priors model in the ambiguous case. Second, we find that with the exception of (α, λ)-pair there is no correlation among any of the model parameters. Third, we find the average level of risk aversion (0.62) to be in line with previous studies (e.g., Harrison and Cox, 2008) and uncorrelated with any of the learning parameters. Next, we investigate these observations in more detail.

We find that spread of the set of priors is large in magnitude (0.66), but the fraction of subjects that are classified as likely to use the multiple priors model is relatively small. So, what is the *extent* of not-Bayesian-ness in our data? To investigate this question, we check whether, for each decision, an MP subject's optimal action could change if they were forced to be Bayesian, depending on the choice of prior from their estimated set of priors. Notice that we only need to check the priors at the endpoints of a subject's set of priors, and therefore, we can re-phrase this question as:

> *For a particular decision, is a subject's optimal action different if we choose the most optimistic prior or the most pessimistic prior?*

We evaluate this question for every decision presented to that subject, and compute the fraction of choices for which we answer "yes" to this question. We weight these fractions by the subject's probability of being MP. The result is presented in Figure 6. The figure shows posterior medians (dots) and 90% credible regions (lines) for each subject around the fraction of decisions that they would reverse based on selecting priors at the endpoints in their set of priors.

Figure 6: "Extent of non-Bayesian-ness"

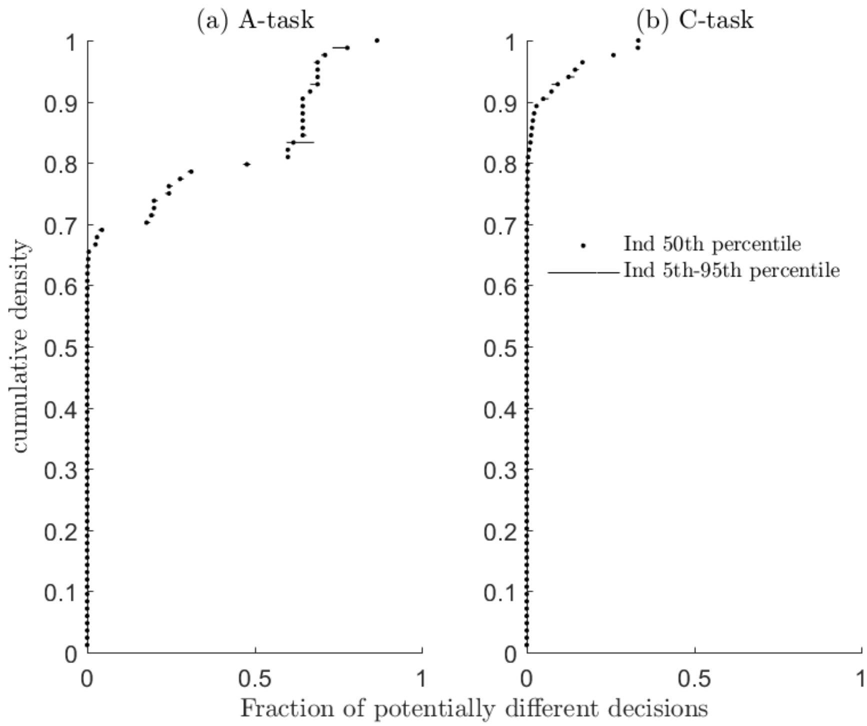

Notes: *A* refers to Ambiguous; *C* refers to Compound. Lines denote the 90% credible regions. B refers to Bayesian, and M refers to multiple priors

Figure 6 shows that MP behavior is more important in the A-task. We can see this in this figure by noting that many more subjects' posterior medians do not fall on the vertical axis in panel (a). In addition, about 20% of subjects were estimated to make different decisions in at least 40% of the A-task, while the corresponding fraction of subjects in the C-task is 0%.

As noted above, we find the precision parameter (λ) is significantly correlated with the learning parameter of the multiple priors model (α). While we believe our treatment of choice precision as a subject-level parameter no different to (say) risk aversion, we note that it is more common in the literature to assume that choice precision is constant across subjects.[7] To investigate the extent to which heterogeneity in λ is important for the conclusion about the prevalence of Bayesian versus Multiple-Priors types, we carry out the same estimation, but with the restriction of the common

[7] Harrison and Rutström (2009), Conte, Hey, and Moffatt (2011), and Harrison, Martínez-Correa, and Swarthout (2015), for example, make this assumption. Ferecatu and Önçüler (2016) on the other hand assumes choice precision is subject-specific.

precision parameter. Table C-5 in the Online Appendix presents the results of the estimation in which we restrict the choice precision to be the same across all subjects. We find that learning estimates do not differ between the two models. This is good news in that the learning parameters and conclusions are robust to the restriction.

Surprisingly, when restricting λ to be the same across subjects we find a substantial reduction in mean risk aversion parameter (from 0.62 to 0.47), and its associated variance estimate. For perspective, a subject with $\gamma_i = 0.62$ would be indifferent between receiving \$1.61 for sure and an equal chance of winning \$10 or nothing, while a subject with $\gamma = 0.47$ would need the sure amount to be \$2.70 to be made indifferent. Figure 7 presents a further investigation of risk aversion in the heterogeneous-λ and homogeneous-λ models. We find that for the heterogeneous-λ model the fraction of risk averse subjects is 0.776 (0.037) and for the homogeneous-λ model the fraction of risk averse subjects drops to 0.688 (0.040). While for our research question, γ was a nuisance parameter, we note that the impact of heterogeneity in λ would have been unknown without our analysis, and *a priori* could have also affected the mixing probabilities and parameters in the multiple priors model. Finally, we note that since we have used the contextual utility model (Wilcox, 2011), this discrepancy is not driven by, for example, payoff differences being uniformly larger for more risk-loving subjects: choice precision, even when utility is normalized, appears to be substantially heterogeneous.

Figure 7: Estimates of Risk Aversion and Precision

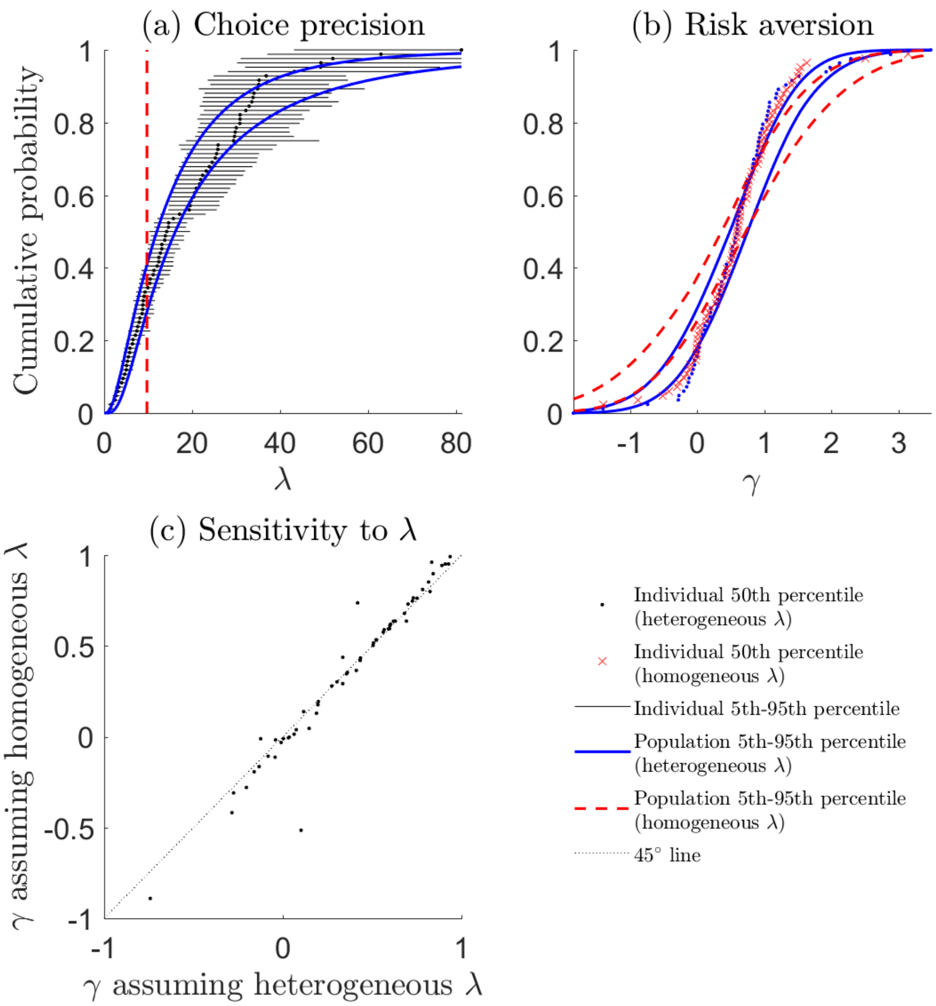

Notes: Each graph describes the distribution of an individual-level parameter. Dots and horizontal black lines show the posterior median and 90% Bayesian credible region for each subject respectively (sorted from lowest to highest median). The thick black line shows a kernel-smoothed density of the posterior medians (normalized so that the maximum density is equal to one). Blue lines show a 90% credible region around the population cumulative density function for the heterogeneous λ model, and red dashed lines show the same credible region estimated from the homogeneous λ model.

5 Conclusion

We ran an economics experiment in order to compare learning under compound risk and under ambiguity using a multiple-priors model of decision making under uncertainty. Participants were required to make sequential choices over pairs of lotteries involving two types of urns: i) a compound urn that was built using a known randomization device, which implied a unique prior; and ii) an ambiguous urn whose composition process was unknown to the participants, and, hence, no unique prior was provided. As successive draws were made from each urn (with replacement), our methodology allowed us to track the best- and worst-case scenarios for the urn composition perceived by the subjects at every drawing round, providing indirect evidence on the set of considered priors at each point in time.

We find that the majority (60%) of subjects learn according to Bayes' rule under both compound risk and under ambiguity, a result that is encouraging for papers that aim to use Bayes' rule to model learning in ambiguous environments (e.g., Bossaerts, Ghirardato, Guarnaschelli, and Zame, 2010). In addition, we find that the second most common type (25%) are subjects that are Bayesian under compound risk, but use a multiple priors model of learning under ambiguity. This result shows that the extent to which behavior under ambiguity differs from behavior under compound risk is relatively moderate. Importantly, we show that restricting subjects' behavior to be consistent with the information provided about the urns (which implies a particular class of priors) leads to incorrect conclusions about learning. Finally, this is one of the first papers that allows for subject-level heterogeneity in preference, learning, and choice precision parameters. We show that while the learning estimates are largely robust, the estimates of risk aversion may be unreliable when restricting choice precision parameters to be the same.

References

Abdellaoui, M., P. Klibanoff, and L. Placido (2015): "Experiments on compound risk in relation to simple risk and to ambiguity," *Management Science*, 61(6), 1306–1322.

Baillon, A., H. Bleichrodt, U. Keskin, O. l'Haridon, and C. Li (2013): "Learning under ambiguity: An experiment using initial public offerings on a stock market," *Economics Working Paper Archive (University of Rennes 1 & University of Caen), Center for Research in Economics and Management (CREM), University of Rennes*, 1.

Bossaerts, P., P. Ghirardato, S. Guarnaschelli, and W. R. Zame (2010): "Ambiguity in asset markets: Theory and experiment," *The Review of Financial Studies*, 23(4), 1325–1359.

Buser, T., L. Gerhards, and J. Van Der Weele (2018): "Responsiveness to feedback as a personal trait," *Journal of Risk and Uncertainty*, 56(2), 165–192.

Chew, S. H., B. Miao, and S. Zhong (2017): "Partial Ambiguity," *Econometrica*, 85(4), 1239–1260.

Cohen, M., I. Gilboa, J.-Y. Jaffray, and D. Schmeidler (2000): "An experimental study of updating ambiguous beliefs," *Risk, Decision and Policy*, 5(2), 123–133.

Conte, A., J. D. Hey, and P. G. Moffatt (2011): "Mixture models of choice under risk," *Journal of Econometrics*, 162(1), 79–88.

Coutts, A. (2019): "Good news and bad news are still news: Experimental evidence on belief updating," *Experimental Economics*, 22(2), 369–395.

Dean, M., and P. Ortoleva (2015): "Is it all connected? A testing ground for unified theories of behavioral economics phenomena," *Working Paper*.

Dominiak, A., P. Dürsch, and J.-P. Lefort (2012): "A dynamic Ellsberg urn experiment," *Games and Economic Behavior*, 75(2), 625–638.

El-Gamal, M. A., and D. M. Grether (1995): "Are people Bayesian? Uncovering behavioral strategies," *Journal of the American statistical Association*, p. 1137–1145.

Ellsberg, D. (1961): "Risk, Ambiguity, and the Savage Axioms," *The Quarterly Journal of Economics*, 75(4), 643–669.

Epstein, L. G., and M. Schneider (2007): "Learning under Ambiguity," *The Review of Economic Studies*, 74(4), 1275–1303.

——— (2010): "Ambiguity and Asset Markets," Working Paper 16181, National Bureau of Economic Research.

Ert, E., and S. T. Trautmann (2014): "Sampling experience reverses preferences for ambiguity," *Journal of Risk and Uncertainty*, 49(1), 31–42.

Ferecatu, A., and A. Önçüler (2016): "Heterogeneous risk and time preferences," *Journal of Risk and Uncertainty*, 53(1), 1–28.

Freeman, D. J., Y. Halevy, and T. Kneeland (2019): "Eliciting risk preferences using choice lists," *Quantitative Economics*, 10(1), 217–237.

Gilboa, I., and D. Schmeidler (1989): "Maxmin expected utility with non-unique prior," *Journal of Mathematical Economics*, 18.

——— (1993): "Updating ambiguous beliefs," *Journal of economic theory*, 59(1), 33–49.

Halevy, Y. (2007): "Ellsberg Revisited: An Experimental Study," *Econometrica*, 75(2), 503–536.

Harrison, G. W., and J. C. Cox (2008): *Risk aversion in experiments*, vol. 12. Emerald Group Pub Ltd.

Harrison, G. W., and E. Elisabet Rutström (2008): "Risk aversion in the laboratory," in *Risk aversion in experiments*, pp. 41–196. Emerald Group Publishing Limited.

Harrison, G. W., J. Martínez-Correa, and J. T. Swarthout (2015): "Reduction of compound lotteries with objective probabilities: Theory and evidence," *Journal of Economic Behavior & Organization*, 119, 32–55.

Harrison, G. W., and E. E. Rutström (2009): "Expected utility theory and prospect theory: One wedding and a decent funeral," *Experimental economics*, 12(2), 133–158.

Holt, C. A., and S. K. Laury (2002): "Risk Aversion and Incentive Effects," *The American Economic Review*, 92(5), 1644–1655.

Kahneman, D., and A. Tversky (1973): "On the psychology of prediction.," *Psychological review*, 80(4), 237.

Knight, F. (1921): *Risk, Uncertainty, and Profit*. University of Chicago Press.

Koop, G., D. J. Poirier, and J. L. Tobias (2007): *Bayesian econometric methods*. Cambridge University Press.

Moreno, O. M., and Y. Rosokha (2016): "Learning under compound risk vs. learning under ambiguity–an experiment," *Journal of Risk and Uncertainty*, 53(2-3), 137–162.

Prokosheva, S. (2016): "Comparing decisions under compound risk and ambiguity: The importance of cognitive skills," *Journal of Behavioral and Experimental Economics*, 64, 94–105.

Qiu, J., and U. Weitzel (2013): "Experimental Evidence on Valuation and Learning with Multiple Priors," SSRN scholarly paper, Social Science Research Network.

——— (2016): "Experimental evidence on valuation with multiple priors," *Journal of Risk and Uncertainty*, 53(1), 55–74.

Savage, L. J. (1954): *The Foundations of Statistics*. Wiley.

Stahl, D. O. (2014): "Heterogeneity of ambiguity preferences," *Review of Economics and Statistics*, 96(4), 609–617.

Wilcox, N. (2011): "Stochastically more risk averse:'A contextual theory of stochastic discrete choice under risk," *Journal of Econometrics*, 162(1), 89–104.

Online Appendix To

Learning Under Uncertainty with Multiple Priors:
Experimental Investigation

James Bland
The University of Toledo

Yaroslav Rosokha
Purdue University

Online Appendix A   Experimental Instructions

Numbered bags hang on top of the blackboard at all times during the experiment. The practice bag has 1 unknown marble in it. (one empty bag for each of the R urns; 4 empty bags for the C urn; one bag with 2 balls of unknown color in case of the first A urn; one bag with 2 balls of unknown color in case of the second A urn)

Experimental Instructions.

Todays experiment will last about 120 minutes. Everyone will earn at least $10. If you follow the instructions carefully you might earn even more money. This money will be paid at the end of the experiment in private and in cash.

It is important that during the experiment you remain SILENT. If you have any questions, or need assistance of any kind, RAISE YOUR HAND but DO NOT SPEAK. One of the experiment administrators will come to you and you may whisper your question to us.

If you talk, laugh, or exclaim out loud, you will be asked to leave and will not be paid. We expect and appreciate your adherence to the instructions.

In total, you will make 180 decisions that affect your potential earnings. Each decision could earn up to $33. At the end of the experiment, **two** of your 180 decisions will be chosen randomly and carried out to determine your actual money earnings. Decisions that will determine your payoff will be selected by rolling dice.

Each decision task will be a set of choices between two lotteries. You can only gain money in these lotteries, you cannot lose any money.

Please click "Continue."

Practice Bag Composition.

We will illustrate decisions and the compensation procedure with the following examples, presented as practice tasks. In these tasks you will choose between two lotteries.

Note, your compensation will not depend on practice decisions. Please direct your attention to the experimenter who will explain the composition of the practice bag.

Press "Continue" once the experimenter finished constructing the bag.

"The outcome of the lottery is based on the color of the ball that will be drawn from the practice bag at the end of practice tasks. In this bag there is one ball of unknown color (either black or white)." Show the bag. Let one of the participants verify that there is one marble in it by touching the bag. "We add two black marbles to the bag and one white marble. The outcome of the first practice task is based on the color of the ball that is drawn from this bag. Please click continue."

Practice Task 1.

For each task you will have two minutes to choose between lotteries labeled A and B. The outcome of the lottery is based on the color of the ball that will be drawn from "practice" bag.

Bag was composed as follows: one ball of unknown color (either black or white), two black balls, and one white ball.

Please choose between lotteries A and B for each question and click 'SUBMIT'.

Option A will pay a fixed amount regardless of the color of the ball. Option B will pay $33 if black ball is drawn, and $5 if white ball is drawn. Please choose between options A and B for each of the questions and press Submit. Notice that you have nine different questions, and for each of them you have to choose A or B. For example question 1): Would you rather have A $6 for sure or B $33 if a black ball is drawn and $5 if a white ball is drawn from the practice bag. Another example is question 9): would you rather have A: $30 for sure or B: $33 if a black ball is drawn and $5 if a white ball is drawn from the practice bag? You need to make separate decision for each of the questions 1–9.

— decision tasks with $33 if black and $5 if white —

Practice Bag Draws.

Let us make three draws from the "practice" bag, replacing the ball after each draw.

Reminder, "practice" bag was composed as follows: one ball of unknown color (either white or black), two black balls, and one white ball.

Ask a participant to draw one ball: "Please draw one ball from this bag." Then ask them to replace the ball into the bag. Repeat two more times. After participants made the three draws (and replaced them into the bag), the experimenter hangs the bag at the top of the blackboard for everyone to see. Next, the experimenter enters draws on own terminal and those records also appear on the participants' screens. "I will record the draws and you can see them on your computer screen."

Practice Task 2.

You will choose between lotteries illustrated below and labeled A and B.

The outcome of the lottery is based on the color of the ball that is drawn from the "practice" bag. Notice that draw will be made from the same "practice" bag and history of draws is available on the right side of this screen. Also notice that the lottery B will have payoffs of either $33 or $5 depending on the color of the ball. But sometimes it will be $33 if black and $5 if white, and other times it will be $33 if white and $5 if black. Each of these will be clearly displayed on your screen.

Reminder, “practice” bag was composed as follows: one ball of unknown color (either black or white), two black balls, and one white ball.

Please choose between lotteries A and B for each question and click ’SUBMIT’. Notice for all pairs of lotteries option B pays **$33 if white ball** is drawn and **$5 if black ball** is drawn.

— decision tasks with $33 if white and $5 if black — + — history box —

Compensation.

Let us demonstrate the compensation procedure. In total you will make 180 decisions, each corresponding to a draw from one of the bags (labeled 1-6). At the end of the experiment you will roll dice to establish decisions that will determine your compensation. At that time you will make a draw from appropriate bags. The procedure will be as follows: you will roll two dice: first die will determine the decision task (for now we will stick with two practice decision tasks), the second die will determine the question selected.

For example, suppose the first die comes up 2 and the second die comes up 3, then the lottery that was randomly chosen is #3 from practice task 2.

Actual Tasks.

Now the tasks for which you will be compensated begin.

In total, you will make 180 decisions that affect your potential earnings (20 tasks with 9 decisions each). Each decision could earn you up to $33. At the end of the experiment, **two of your 180 decisions** will be chosen randomly and carried out to determine your actual money earnings. The decision that will determine your payoff will be selected by rolling two dice. The first die will determine the task number, the second die will determine the decision number. This procedure will be repeated twice, so the two decisions selected will be independent.

Bag 1 Composition.

For tasks 1 and 2 the outcome of the lotteries is based on the color of the ball that will be drawn from bag 1. At this time please direct your attention to the experimenter who will explain the composition of bag 1.

Please press “Continue” once the experimenter finished constructing the bag.

“This bag is empty” Show the bag.” Let one of the participants verify that there are no balls there by touching the bag. “We add two balls and two white balls to bag 1.” Place the bag at the top of the blackboard so that everyone can see it for the duration of the experiment.

Task 1.

The outcome of the lottery is based on the color of the ball that will be drawn from bag 1. Reminder, bag 1 was composed as follows: two black balls, and two white balls.

Please choose between lotteries A and B for each question and click 'SUBMIT'. Notice that for all pairs of lotteries option B pays **$33 if black ball** is drawn and **$5 if white ball** is drawn.

— decision tasks with $33 if black and $5 if white —

Task 2.

The outcome of the lottery is based on the color of the ball that will be drawn from bag 1. Reminder, bag 1 was composed as follows: two black balls, and two white balls.

Please choose between lotteries A and B for each question and click 'SUBMIT'. Notice that for all pairs of lotteries option B pays **$33 if white ball** is drawn and **$5 if black ball** is drawn.

— decision tasks with $33 if white and $5 if black —

Tasks 3 and 4 follow the same procedure. Specifically, Tasks 3 corresponds to Bag 2 which contained 3 black balls and 1 white ball and Task 4 corresponds to Bag 3 which contained 3 white balls and 1 black ball. Task 3 involves decisions with $33 if black and $5 if white;

Bag 4 Composition.

For tasks 5 through 10 the outcome of the lotteries is based on the color of the ball that will be drawn from bag 4. At this time please direct your attention to the experimenter who will explain the composition of bag 4.

Please press continue once the experimenter finished constructing the bag.

"These four bags are empty." Show the bags, let one of the participants verify that there are no balls there by touching them. Show them a box and let one of the students verify that it is empty. "We add one black and three white balls to the first bag and place it in the box. We add two black and two white balls to the second bag and place it in the box. We add two black and two white balls to the third bag and place it in the box. We add three black and one white ball to the fourth bag and place it in the box." Shake the box and let one of the participants draw one bag from the box. "That is Bag 4 used for tasks 5 through 10."

Task 5.

The outcome of the lottery is based on the color of the ball that is drawn from bag 4.

Reminder, bag 4 was composed as follows: 1/4 chance of one white and three black balls; 2/4 chance of 2 white and 2 black balls; 1/4 chance of three white and one black balls.

Please choose between lotteries A and B for each question and click 'SUBMIT'. Notice that for all pairs of lotteries option B pays **$33 if black ball** is drawn and **$5 if white ball** is drawn.

— decision tasks with $33 if white and $5 if black —

Task 6 is the same as Task 5 with the exception that the decision tasks involve $33 if white and $5 if black.

Bag 4 Draws.

Reminder, bag 4 was composed as follows: 1/4 chance of one white and three black balls; 2/4 chance of 2 white and 2 black balls; 1/4 chance of three white and one black balls.

Now, let us make three draws from bag 4, replacing the ball after each draw.

Ask a participant to draw one ball. Then ask them to replace the ball into the bag. Repeat two more times. After participants made the three draws (and replaced them into the bag), the experimenter hangs the bag at the top of the blackboard for everyone to see. Next, the experimenter enters draws on own terminal. "I will record the draws and you can see them on your computer screen".

Task 7.

The outcome of the lottery is based on the color of the ball that is drawn from bag 4.

Reminder, bag 4 was composed as follows: 1/4 chance of one white and three black balls; 2/4 chance of 2 white and 2 black balls; 1/4 chance of three white and one black balls.

Please choose between lotteries A and B for each question and click 'SUBMIT'. Notice that for all pairs of lotteries option B pays **$33 if black ball** is drawn and **$5 if white ball** is drawn.

— decision tasks with $33 if white and $5 if black — + — history box —

Task 8 is the same as Task 7 with the exception that the decision tasks involve $33 if white and $5 if black.

Bag 4 Draws.

Reminder, bag 4 was composed as follows: 1/4 chance of one white and three black balls; 2/4 chance of 2 white and 2 black balls; 1/4 chance of three white and one black balls.

Now, let us make three draws from bag 4, replacing the ball after each draw.

Ask a participant to draw one ball. Then ask them to replace the ball into the bag. Repeat two more times. After participants made the three draws (and replaced them into the bag), the experimenter hangs the bag at the top of the blackboard for everyone to see. Next, the experimenter enters draws on own terminal.

Tasks 9 and **10** are the same as Tasks 7 and 8, with an updated history of draws.

Bag 5 Composition.

For tasks 11 through 16 the outcome of the lotteries is based on the color of the ball that will be drawn from bag . At this time please direct your attention to the experimenter who will explain the composition of bag 5.

Please press continue once the experimenter finished constructing the bag.

"In this bag there are two balls of unknown color (each could be either black or white)." Show the bag. Let one of the participants verify that there are two balls in the bag by touching it. "We add one black ball to the bag and one white ball to the bag. Please click Continue."

Tasks 11 through 16 are the same as Tasks 5 through 10 with the difference that the bag is the Ambiguous bag. Note that in half of the treatments the Ambiguous bag was presented before the Compound bag.

Tasks 17 through 20 are the same format at Tasks 11 through 16 with the difference that a different bag (**Bag 6**) is used. The construction process of Bag 6 is the same as Bag 5. Another difference is that after the first and the second set of draws subjects face only one task ($33 if black and $5 if white for Task 19, and $33 if white and $5 if black for Task 20).

Figure A-1 presents the screenshot of one of the Tasks from the experiment. We highlight three component of the screen: (1) Information regarding which bag is used and the reminder of that bag's composition; (2) Decision tasks; (3) History of draws and the draw summary (if any). Note that although the figure presents the decisions with $33 if black and $5 if white, half of the decisions tasks involved decisions with option B being $33 if white and $5 if black.

Figure A-1: Screenshot of the Experimental Interface.

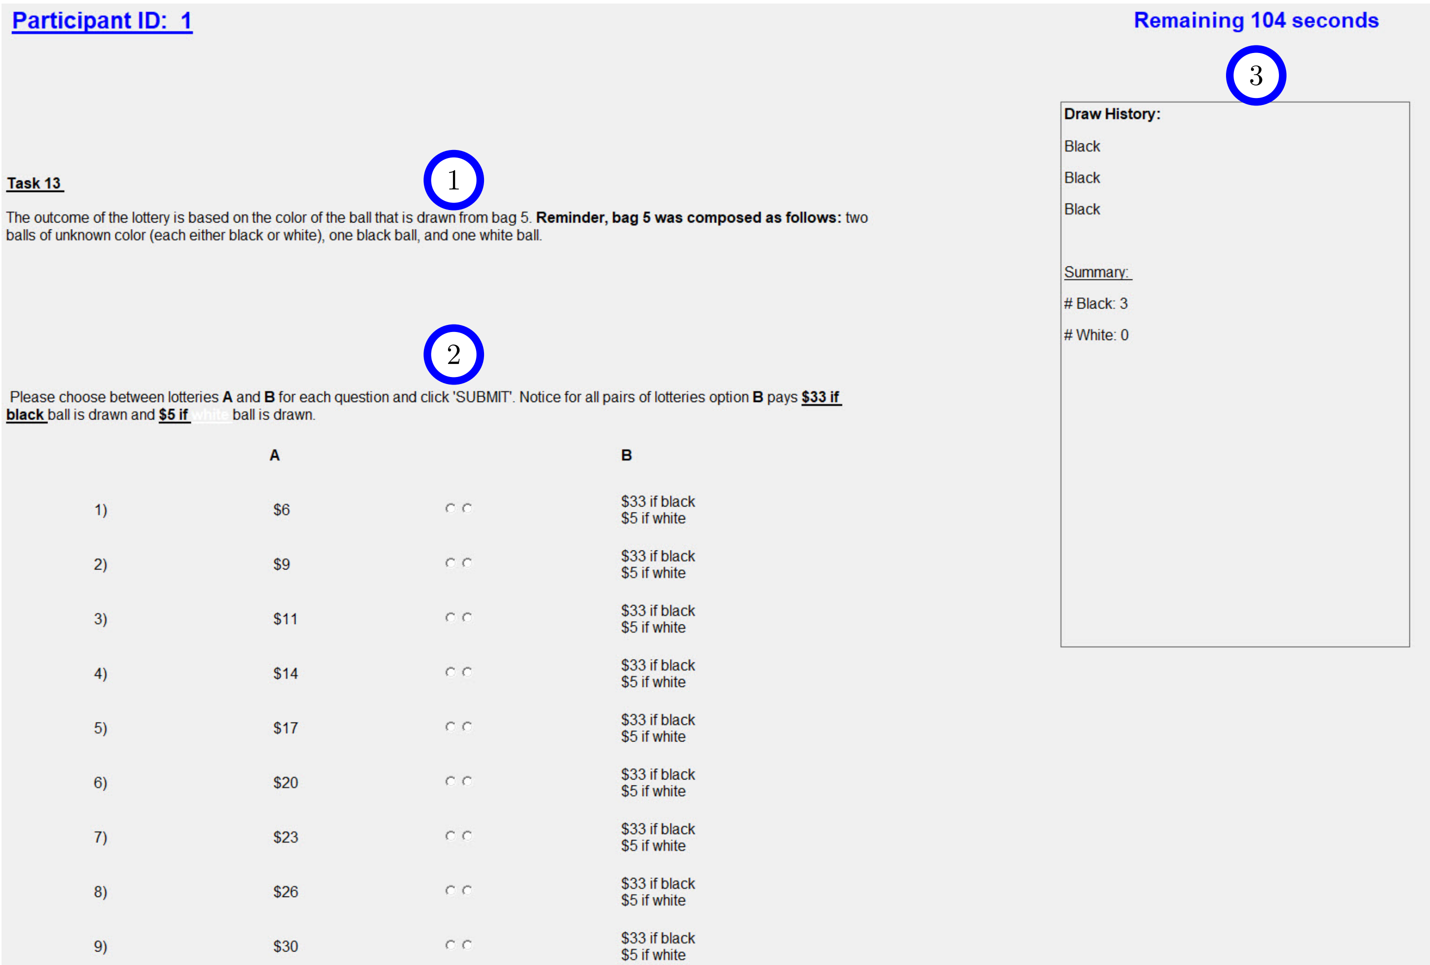

Notes: (1) Information regarding which bag is used and the reminder of that bag's composition; (2) Decision tasks; (3) History of draws and the draw summary (if any).

Online Appendix B Robustness Checks

Table B-1: Summary of mixing probability estimates

(a) With Data from Moreno & Rosokha (2016)

Mixing probabilities – joint		
AB CB	0.510	(0.055)
AB CM	0.094	(0.040)
AM CB	0.229	(0.053)
AM CM	0.166	(0.045)
Mixing probabilities – marginal		
CM	0.260	(0.056)
AM	0.396	(0.054)
Pr(CM > AM)	0.034	
Prob modal type		
AB CB	0.998	
AB CM	0.000	
AM CB	0.002	
AM CM	0.000	

Marginal log-likelihood: -6703 (mean)

(b) Restricting Beta priors to [.25,.75]

Mixing probabilities – joint		
AB CB	0.602	(0.063)
AB CM	0.035	(0.028)
AM CB	0.293	(0.059)
AM CM	0.069	(0.032)
Mixing probabilities – marginal		
CM	0.104	(0.040)
AM	0.362	(0.061)
Pr(CM > AM)	0.000	
Prob modal type		
AB CB	0.996	
AB CM	0.000	
AM CB	0.004	
AM CM	0.000	

Marginal log-likelihood: -3630 (mean)

Notes: Panel (a) shows estimates for Beta priors, including data from both experiments. Panel (b) shows estimates for restricted Beta priors. Reported mixing probabilities are posterior means (standard deviations). "Prob modal type" reports the posterior probability that each type is the most prevalent in the population. AB-CB, AB-CM, AM-CB, and AM-CM are the four possible types that subjects could be classified into from the set {Bayesian, Multiple priors} $\times$ {A-task, C-task}.

Table B-2: Unrestricted beta priors parameter estimates, using additional data from Moreno & Rosokha (2016)

	Common		Compound - Bayes		Ambiguous - Bayes		Ambiguous - Multiple Priors			
	γ	$\log\lambda$	$\log N_0$	p_0	$\log N_0$	p_0	α	$\log N_0$	p_{0mid}	p_{0sp}
Mean										
	0.65	3.17	1.08	0.48	0.85	0.49	0.17	2.38	0.52	0.20
	(0.12)*	$(0.16)^a$	(0.14)*	$(0.01)^a$	(0.19)*	$(0.01)^a$	$(0.10)^a$	(0.42)*	$(0.02)^a$	$(0.07)^a$
Variance & Correlation										
γ	2.46^a									
$\log\lambda$	0.63	5.06^a								
$\log N_0$	0.02	0.05	0.81^a							
p_0	-0.01	0.02	0.08	0.08^a						
$\log N_0$	-0.62	-0.76*	0.21	-0.02	8.91^a					
p_0	0.03	0.06	0.22	0.05	0.00	0.19^a				
α	-0.48	-0.65*	0.09	-0.04	0.68*	-0.00	0.29^a			
$\log N_0$	-0.12	-0.05	-0.06	-0.00	0.36*	-0.05	0.12	0.88^a		
p_{0mid}	-0.07	0.02	0.38	0.07	0.09	0.24	0.04	-0.07	0.13^a	
p_{0sp}	-0.03	-0.17	-0.35	-0.03	-0.01	-0.14	0.01	0.10	-0.15	0.14^a

Notes: Table shows posterior means (standard deviations). * indicates that a 95 percent Bayesian credible region does not include zero. ^{a} indicates that stars are suppressed because these parameters can only be positive. γ - risk-aversion. λ - logit choice precision. N_0 - prior strength. p_0 prior mean. α - likelihood for discarding priors. p_{0mid} midpoint of the set of prior means. p_{0sp} spread of the set of prior means.

Table B-3: Parameter estimates, restricting Beta priors to [.25,.75]

	Common		Compound - Bayes		Ambiguous - Bayes		Ambiguous - Multiple Priors			
	γ	$\log\lambda$	$\log N_0$	p_0	$\log N_0$	p_0	α	$\log N_0$	p_{0mid}	p_{0sp}
Mean										
	0.56	2.58	-3.92	0.02	-8.41	0.52	0.88	2.30	0.00	0.00
	(0.11)*	(0.11)^a	(0.80)*	(0.06)^a	(1.13)*	(0.02)^a	(0.09)^a	(0.22)*	(0.00)^a	(0.00)^a
Variance & Correlation										
γ	5.63^a									
$\log\lambda$	0.23	4.36^a								
$\log N_0$	-0.14	-0.06	0.45^a							
p_0	0.38*	0.28	0.07	0.02^a						
$\log N_0$	-0.54	-0.41	0.43*	-0.28*	11.06^a					
p_0	-0.64	0.07	0.10	-0.24*	0.36	0.49^a				
α	-0.01	0.43	0.12	0.11*	-0.02	0.21*	0.24^a			
$\log N_0$	-0.09	-0.16	0.42*	-0.07	0.29*	0.07	-0.00	2.23^a		
p_{0mid}	0.06	-0.01	-0.05	0.02	-0.07*	-0.05*	-0.08	-0.11*	0.00^a	
p_{0sp}	0.05	-0.01	-0.05	0.01	-0.07*	-0.05*	-0.08	-0.12*	0.97*	0.00^a

Notes: Table shows posterior means (standard deviations). * indicates that a 95 percent Bayesian credible region does not include zero. ^{a} indicates that stars are suppressed because these parameters can only be positive. γ - risk-aversion. λ - logit choice precision. N_0 - prior strength. p_0 prior mean. α - likelihood for discarding priors. p_{0mid} midpoint of the set of prior means. p_{0sp} spread of the set of prior means.

Online Appendix C Additional Tables and Figures

Figure C-2: Task Order

Stage	1	2	3	4
Treatment 1	R	C	A	A
Treatment 2	R	A	C	A

Figure C-3: Individual posterior type probabilities.

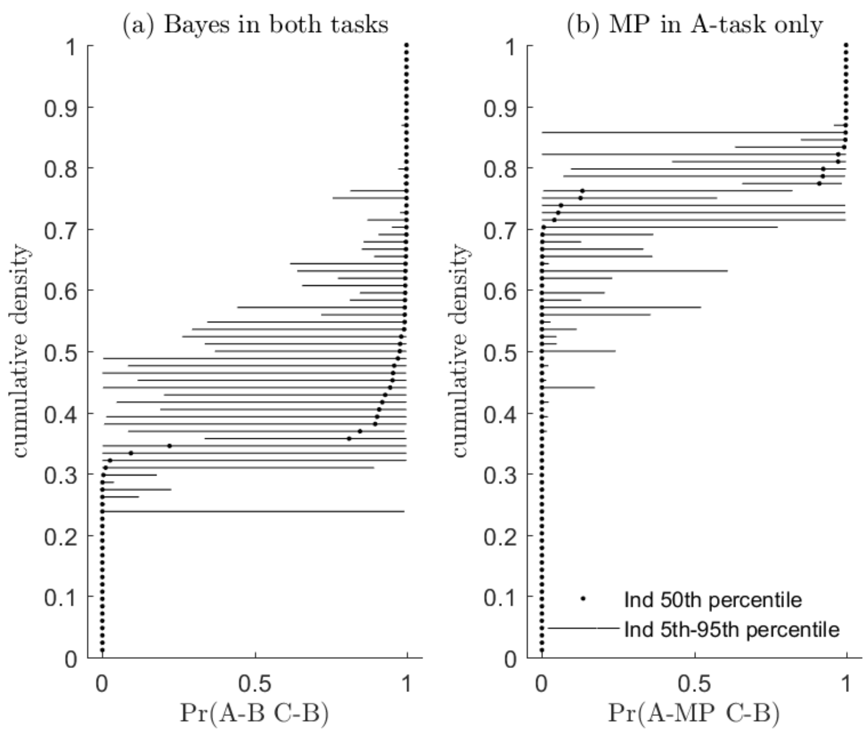

Notes: Panel (a) shows individual posterior probabilities that a subject is Bayesian in both tasks. Panel (b) shows individual posterior probabilities that a subject is Bayesian in the compound task, but Multiple Priors in the ambiguous task. Dots (lines) show posterior medians (90% credible regions) for these probabilities for each subject.

While Table 1 shows estimates of the fraction of each type in the population, we can also assign posterior probabilities to each subject being each type. We do this for the two most prevalent types, shown in Figure C-3, where dots show a subject's posterior median probability of being that type, and the lines show a 90% credible region around that. For a reasonable fraction of subjects the credible region is quite small. If a subject's credible region covers only probabilities close to 1, we can be very certain that that subject is that type. On the other hand, if the credible region is bunched up around zero, the we can be very sure that the subject is *not* that type. Wide credible regions in panel (a) of this Figure mostly correspond to wide credible regions in panel (b): our uncertainty about these subjects is whether they are B or MP in the A task.

Table C-4: Restricted λ beta priors parameter estimates

	Common	Compound - Bayes		Ambiguous - Bayes		Ambiguous - Multiple Priors			
	γ	$\log N_0$	p_0	$\log N_0$	p_0	α	$\log N_0$	p_{0mid}	p_{0sp}
Mean									
	0.47	1.40	0.49	1.23	0.50	0.83	5.70	0.50	0.54
	(0.12)*	(0.20)*	$(0.01)^a$	(0.27)*	$(0.01)^a$	$(0.12)^a$	(0.48)*	$(0.03)^a$	$(0.09)^a$
Variance & Correlation									
γ	3.15^a								
$\log N_0$	-0.09	0.77^a							
p_0	0.01	-0.04	0.07^a						
$\log N_0$	-0.39	0.14	0.00	16.54^a					
p_0	-0.08	0.28	-0.02	0.02	0.23^a				
α	0.28	-0.08	-0.02	-0.67*	0.00	0.29^a			
$\log N_0$	-0.11	0.30	-0.04	-0.10	0.22	0.15	1.42^a		
p_{0mid}	0.05	-0.06	0.02	-0.02	-0.02	0.01	-0.06	0.08^a	
p_{0sp}	-0.26	0.55	-0.01	-0.00	0.23	-0.01	0.24	-0.05	0.23^a
$\log \lambda$	2.28	(0.02)							

Notes: Table shows posterior means (standard deviations). * indicates that a 95 percent Bayesian credible region does not include zero. [a] indicates that stars are suppressed because these parameters can only be positive. γ - risk-aversion. λ - logit choice precision. N_0 - prior strength. p_0 prior mean. α - likelihood for discarding priors. p_{0mid} midpoint of the set of prior means. p_{0sp} spread of the set of prior means.

Table C-5: Estimates from the restricted Beta priors model, setting choice precision λ_i to be constant across subjects.

(a) Estimates summary

	γ	α	$\log N_0$	p_{0mid}	p_{0sp}
Mean	0.47	0.83	5.70	0.50	0.54
	(0.12)*	(0.12)*	(0.48)*	(0.03)*	(0.09)*
Variance / correlation					
γ	12.37	-	-	-	-
	$(14.20)^a$				
α	0.28	0.10	-	-	-
	(0.42)	$(0.05)^a$			
$\log N_0$	-0.11	0.15	2.14	-	-
	(0.22)	(0.14)	$(1.15)^a$		
p_{0mid}	0.05	0.01	-0.06	0.01	-
	(0.20)	(0.14)	(0.25)	$(0.00)^a$	
p_{0sp}	-0.26	-0.01	0.24	-0.05	0.06
	(0.21)	(0.15)	(0.26)	(0.48)	$(0.03)^a$
$\log \lambda$	2.28	(0.02)			

(b) Mixing probabilities

Mixing probabilities – joint		
AB CB	0.602	(0.063)
AB CM	0.035	(0.028)
AM CB	0.293	(0.059)
AM CM	0.069	(0.032)
Mixing probabilities – marginal		
CM	0.104	(0.040)
AM	0.362	(0.061)
Pr(CM > AM)	0.000	
Prob modal type		
AB CB	0.996	
AB CM	0.000	
AM CB	0.004	
AM CM	0.000	

Notes: Table shows posterior means (standard deviations). * indicates that a 95 percent Bayesian credible region does not include zero. ^{a} indicates that stars are suppressed because these parameters can only be positive. γ - risk-aversion. λ - logit choice precision. N_0 - prior strength. p_0 prior mean. α - likelihood for discarding priors. p_{0mid} midpoint of the set of prior means. p_{0sp} spread of the set of prior means. AB-CB, AB-CM, AM-CB, and AM-CM are the four possible types that subjects could be classified into from the set $\{\text{Bayesian, Multiple priors}\} \times \{\text{A-task, C-task}\}$

Online Appendix D Notes on Bayesian estimation techniques

Table D-6: List of individual parameters used in Beta priors specifications

	Description	Initial value / prior mean
	Parameters common to all types	
1	γ, CRRA utility function parameter	0.5
2	$\log(\lambda)$, logistic choice precision	$2 \implies \lambda \approx 7.4$
	A-B type	
3	$\log(N_0^{A-B})$ Strength of prior	log(2), i.e. uniform
4	$\Phi^{-1}(p_0^{A-B})$, mean of prior	0, i.e. uniform
	A-MP type	
5	$\Phi^{-1}(\alpha)$, extent discarding unlikely priors	0
6	$\log(N_0^{A-MP})$ Strength of set of priors	log(2)
7/8	Parameters governing endpoints of set of priors: $\underline{p}_0 = \Phi(\theta_7) - \Phi(\theta_7)\Phi(\theta_8)$, $\overline{p}_0 = \Phi(\theta_7) + (1 - \Phi(\theta_7))\Phi(\theta_8)$	$\theta_7 = \theta_8 = 0$
	C-B type	
9	$\log(N_0^{A-B})$ Strength of prior	log(2), i.e. uniform
10	$\Phi^{-1}(p_0^{A-B})$, mean of prior	0, i.e. uniform
	C-MP type	
11	$\Phi^{-1}(\alpha)$, extent discarding unlikely priors	0
12	$\log(N_0^{A-MP})$ Strength of set of priors	log(2)
13/14	Parameters governing endpoints of set of priors: $\underline{p}_0 = \Phi(\theta_7) - \Phi(\theta_7)\Phi(\theta_8)$, $\overline{p}_0 = \Phi(\theta_7) + (1 - \Phi(\theta_7))\Phi(\theta_8)$	$\theta_7 = \theta_8 = 0$

We assume that subjects behave according to exactly one of four models of decision-making. These models are:

- (A-B C-B) Bayesian with subjective priors (henceforth B) in both the A- and C-tasks, indexed by $\tau = 1$
- (A-B C-MP) Bayesian in the A-task, Epstein and Schneider (2007) multiple priors (henceforth MP) in the C-task, indexed by $\tau = 2$
- (A-MP C-B) MP in the A-task, B in the C-task, indexed by $\tau = 3$
- (A-MP C-MP) MP in both tasks, indexed by $\tau = 4$

Table D-7: List of individual parameters used in Simplex priors specifications

	Description	Initial value / Prior mean
	PARAMETERS COMMON TO ALL TYPES	
1	γ, CRRA utility function parameter	0.5
2	$\log(\lambda)$, logistic choice precision	$2 \implies \lambda \approx 7.4$
	A-B TYPE	
3	$\Phi(\theta_3)$ = prior that the urn has 1 black marble, conditional on it not having 2	0
4	$\Phi(\theta_4)$ = prior that the urn has 2 black marbles	0
	A-MP TYPE	
5	$\theta_5 = \Phi^{-1}(\alpha)$, extent discarding unlikely priors	0
6	$\Phi(\theta_6)$ = smallest prior probability assigned to urn containing 2 black marbles	-2
7/8	Parameters governing the minimum probabilities assigned to there being 1 or 3 black marbles. $\underline{p}_1 = (1-\Phi(\theta_6))\Phi(\theta_7)\Phi(\theta_8)$, $\underline{p}_1 = (1-\Phi(\theta_6))\Phi(\theta_7)(1-\Phi(\theta_8))$	$\theta_7 = \theta_8 = -2$
	C-B TYPE	
9	$\Phi(\theta_9)$ = prior that the urn has 1 black marble, conditional on it not having 2	0
10	$\Phi(\theta_{10})$ = prior that the urn has 2 black marbles	0
	C-MP TYPE	
11	$\theta_{11} = \Phi^{-1}(\alpha)$, extent discarding unlikely priors	0
12	$\Phi(\theta_{12})$ = smallest prior probability assigned to urn containing 2 black marbles	-2
13/14	Parameters governing the minimum probabilities assigned to there being 1 or 3 black marbles. $\underline{p}_1 = (1-\Phi(\theta_{12}))\Phi(\theta_{13})\Phi(\theta_{14})$, $\underline{p}_1 = (1-\Phi(\theta_{12}))\Phi(\theta_{13})(1-\Phi(\theta_{14}))$	$\theta_{13} = \theta_{14} = -2$

Each model specifies a likelihood function mapping individual-level parameters θ_i into a probability distribution over actions Y_i. We denote these likelihood functions as:

$$p(Y_i \mid \theta_i, \tau = 1),\ p(Y_i \mid \theta_i, \tau = 2),\ p(Y_i \mid \theta_i, \tau = 3),,\ p(Y_i \mid \theta_i, \tau = 4)$$

We assume that subjects' behavior is independent, so conditional on knowing all subjects behavioral parameters θ, and their types τ, we can construct the likelihood of observing *all* subjects' data as:

$$p(Y \mid \theta, \tau) = \prod_{i=1}^{N} p(Y_i \mid \theta_i, \tau_i)$$

We aim to simulate the posterior distribution $p(\beta, \Sigma, \rho \mid Y)$. β and Σ govern the distribution of θ, and ρ governs the distribution of τ. To this end, we augment the data with the individual-level parameters θ and τ to get the joint posterior distribution of $(\beta, \Sigma, \rho, \theta, \tau)$:

$$\begin{aligned}
p(\beta, \Sigma, \rho, \theta, \tau) &\propto p(Y \mid \beta, \Sigma, \rho, \theta, \tau) p(\beta, \Sigma, \rho, \theta, \tau) && (12)\\
&= \prod_{i=1}^{N} \left[p(Y_i \mid \beta, \Sigma, \rho, \theta, \tau\right] p(\beta, \Sigma, \rho, \theta, \tau) && (13)\\
&= \prod_{i=1}^{N} \left[\sum_{\tau} p(Y_i \mid \theta_i, \tau) I(\tau_i = \tau) \rho_\tau\right] p(\theta \mid \beta, \Sigma) p(\beta, \Sigma, \rho) && (14)\\
&= \prod_{i=1}^{N} \left[\sum_{\tau} p(Y_i \mid \theta_i, \tau) I(\tau_i = \tau) \rho_\tau\right] p(\theta \mid \beta, \Sigma) p(\beta, \Sigma) p(\rho) && (15)
\end{aligned}$$

where $p(Y_i \mid \theta_i, \tau)$ is subject i's likelihood conditional on having parameters θ_i and being type τ. The final equality assumes that for the prior distribution, (β, Σ) is independent of ρ.

Using Gibbs sampling, we can draw from this distribution if we can draw from its conditionals. Broadly, this will be done in five steps (including an initialization):

0. Initialization: Choose initial values. These are summarized in Tables D-6 and D-7 for the Beta and Simplex priors specifications respectively.

1. Draw from $\beta, \Sigma \mid \theta, \rho, \tau, Y$. Inspection of (15) yields that:

$$p(\beta, \Sigma \mid \theta, \rho, \tau, Y) \propto p(\theta \mid \beta, \Sigma) p(\beta, \Sigma, \rho) \tag{16}$$

   if we use a Normal-Inverse-Wishart prior $(\beta, \Sigma) \sim NIW(\underline{M}, \underline{L}, \underline{P}, \underline{V})$, with (β, Σ) independent of ρ in the prior distribution, then:

$$\beta, \Sigma \mid \theta, \rho, \tau, Y \sim \mathcal{NIW}(\overline{M}, \overline{L}, \overline{P}, \overline{V}) \tag{17}$$

$$\overline{M} = \frac{\underline{L}\underline{M} + N\bar{\theta}}{\underline{L} + N} \tag{18}$$

$$\overline{L} = \underline{L} + N \tag{19}$$

$$\overline{V} = \underline{V} + N \tag{20}$$

$$\overline{P} = \underline{P} + \sum_{i=1}^{N} (\theta_i - \bar{\theta})'(\theta_i - \bar{\theta}) + \frac{\underline{L}N}{\underline{L} + N} (\bar{\theta} - \underline{M})'(\bar{\theta} - \underline{M}) \tag{21}$$

where $\bar{\theta} = \frac{1}{N} \sum_{i=1}^{N} \theta_i$. See Koop, Poirier, and Tobias (2007, Ex. 12.1) for a more general derivation of this result. We can therefore draw from $\beta, \Sigma \mid \theta, \rho, \tau, Y$ as follows:

(a) Draw $\Sigma \mid \theta, Y \sim IW(\overline{P}, \overline{V})$

(b) Draw $\beta \mid \Sigma, \theta, Y \sim N\left(\overline{M}, \Sigma/\overline{L}\right)$

We set the prior mean vector $\underline{M}$ equal to our starting values of θ (See Tables D-6 and D-7.), $\underline{P}$ equal to the identity matrix, $\underline{T} = 1$, and $\underline{V}$ equal to the number of elements in β plus 2.[8] Note that by choosing small $\underline{L}$ and $\underline{V}$, the (conditional) posterior of (β, Σ) is driven largely by θ, our estimates of the individual-level parameters.

2. Draw $\theta \mid \beta, \Sigma, \rho, \tau = \tau_k, Y$ for each model τ_k. The relevant component of (15) is:

$$p(\theta_i \mid \theta_{-i}\beta, \Sigma, \rho, \tau = \tau_k, Y) \propto p(Y_i \mid \theta_i, \tau = \tau_k) p(\theta_i \mid \beta, \Sigma) \quad \forall i \tag{22}$$

As $p(Y_i \mid \theta_i, \tau = \tau_k)$ is typically non-standard, we use a Metropolis-Hastings algorithm to perform this step.

3. Draw $\tau \mid \beta, \Sigma, \rho, \theta, Y$, and update θ to be the one from above specific to this draw. The relevant component of (15) is:

$$p(\tau_{i,k} \mid \beta, \Sigma, \rho, \theta, Y) \propto p(Y_i \mid \theta_i, \tau_{i,k}) \rho_k \tag{23}$$

$$\implies p(\tau_{i,k} \mid \beta, \Sigma, \rho, \theta, Y) = \frac{p(Y_i \mid \theta_i, \tau_{i,k}) \rho_k}{\sum_l p(Y_i \mid \theta_i, \tau_{i,l}) \rho_l} \tag{24}$$

Note that the simulated values of (24) can be used to assign posterior probabilities to individual subjects being each type. While we do not need to store these to make statements about the posterior moments of the population-level parameters (β, Σ, ρ), we may nonetheless wish to store these if we want to say things about specific subjects.

[8] For the mean of this prior distribution to exist, $\underline{V}$ must be at least the number of elements in β plus 1.

4. Draw $\rho \mid \beta, \Sigma, \theta, \tau, Y$. From (15):

$$p(\rho \mid \beta, \Sigma, \theta, \tau, Y) \propto p(\rho) \prod_{i=1}^{N} \rho_{\tau_i} \tag{25}$$

$$= p(\rho) \prod_{k} \rho_k^{\sum_i I(\tau_i = k)} \tag{26}$$

If we assume a Dirichlet prior:

$$p(\rho) \propto \prod_{k} \rho_k^{\underline{\alpha}_k} \tag{27}$$

then:

$$p(\rho \mid \beta, \Sigma, \theta, \tau, Y) \propto \prod_{k} \rho_k^{\underline{\alpha}_k} \prod_{k} \rho_k^{\sum_i I(\tau_i = k)} \tag{28}$$

$$= \prod_{k} \rho_k^{\underline{\alpha}_k + \sum_i I(\tau_i = k)} \tag{29}$$

$$\rho \mid \beta, \Sigma, \theta, \tau, Y \sim \text{Dirichlet}(\overline{\alpha}_1, \overline{\alpha}_2, \ldots, \overline{\alpha}_K) \tag{30}$$

$$\overline{\alpha}_k = \underline{\alpha}_k + \sum_i I(\tau_i = k) \tag{31}$$

We use $\underline{\alpha}_1 = \underline{\alpha}_2 = \underline{\alpha}_3 = \underline{\alpha}_4 = 1$, which implies the following about the prior distribution:

- The prior means of each mixing probability are all equal to 25%. (i.e. Each of the four types are equally prevalent in expectation)
- The marginal distribution of types in each tasks are:

$$p(\text{subject } i \text{ is type A-MP}) \sim iid\text{Beta}(2, 2) \tag{32}$$

$$p(\text{subject } i \text{ is type C-MP}) \sim iid\text{Beta}(2, 2) \tag{33}$$