Blame It on the Coin Flip: Preferences for Randomization and Regret

Online Appendix: Omitted Proofs

Johannes C. Hoelzemann *and* Moritz Loewenfeld

November 22, 2025

Appendix A: Proofs

A.1 Proof of Proposition 1

Proof. As the first claim is obvious, we proceed showing the second claim. Let $R(\cdot)$ be strictly increasing. Fix two lotteries A and B, and any $\lambda \in (0, 1]$. Change the payoffs of lottery A such that $x_k^{A,\pi} = x_k^A + \pi$, $x_{\omega \neq k}^{A,\pi} = x_\omega^A$. That is, the payoff in state k is changed by π, while the other payoffs are kept constant.

Moreover, let

$$V(\gamma, \pi) = \sum_{\omega \in \Omega} p_\omega \, R\Big[(2\gamma - 1)\Big(\lambda \Delta_\omega(\pi) + (1 - \lambda)\Delta_{EU}(\pi)\Big)\Big], \tag{A.1.1}$$

$$\gamma^*(\pi) = \arg\max_{\gamma \in [0,1]} V(\gamma, \pi). \tag{A.1.2}$$

Then, the claim is that

$$\pi_1 < \pi_2 \quad \Longrightarrow \quad \gamma^*(\pi_1) \;\leq\; \gamma^*(\pi_2). \tag{A.1.3}$$

We use monotone comparative-statics results to prove the claim. Applying

Theorems 5 and 6 in Milgrom and Shannon (1994), the claim is true if the cross-derivative is positive.[15] For each state ω, we write $D_\omega(\pi) = \lambda\Delta_\omega(\pi) + (1 - \lambda)\Delta_{EU}(\pi)$. The cross-derivative is given by:

$$\frac{\partial^2 V}{\partial\gamma\,\partial\pi} = 2\lambda p_k \frac{\partial D_k}{\partial\pi}\Big[R'\Big((2\gamma-1)D_k\Big) + (2\gamma-1)D_k\,R''\Big((2\gamma-1)D_k\Big)\Big]. \quad \text{(A.1.4)}$$

Because Δ_k and $\Delta_{EU}(\pi)$ are strictly increasing in π, $\frac{\partial D_k}{\partial\pi} > 0$. Therefore, expression (A.1.4) is positive if the term in square brackets is positive. This is the case if $R'(z) + zR''(z) \geq 0$ for any $z \in \mathbb{R}$, which proves claim 2.

To establish claim 3, we first require minimal additional notation. Consider two lottery pairs (A, B) and $(A^\dagger, B^\dagger)$, with random variables $F = \{(p_\omega, \Delta_\omega)\}_{\omega\in\Omega}$ and $F^\dagger = \{(p^\dagger_\omega, \Delta^\dagger_\omega)\}_{\omega\in\Omega}$, such that $F^\dagger$ is a mean-preserving spread of F. Denote $D_\omega = \lambda\Delta_\omega + (1-\lambda)\Delta_{EU}$ and $D^\dagger_\omega = \lambda\Delta^\dagger_\omega + (1-\lambda)\Delta_{EU}$. We define the objective functions as follows:

$$V(\gamma) := \sum_\omega p_\omega R\Big((2\gamma-1)D_\omega\Big), \quad V^\dagger(\gamma) := \sum_\omega p_\omega R\Big((2\gamma-1)D^\dagger_\omega\Big), \quad \text{(A.1.5)}$$

and let $\gamma^* = \arg\max_{\gamma\in[0,1]} V(\gamma)$, $\gamma^{\dagger*} = \arg\max_{\gamma\in[0,1]} V^\dagger(\gamma)$.

Our goal is to find conditions under which $0.5 < \gamma^{\dagger*} \leq \gamma^*$ as $\Delta_{EU} > 0$. We take the derivative with respect to γ:

$$V'(\gamma) = 2\sum_\omega p_\omega R'\Big((2\gamma-1)D_\omega\Big)D_\omega. \quad \text{(A.1.6)}$$

Evaluated at $\gamma = 0.5$, this expression simplifies to

$$V'(\gamma = 0.5) = 2\sum_\omega p_\omega R'\Big(0\Big)\Delta_{EU} = 2R'\Big(0\Big)\Delta_{EU}, \quad \text{(A.1.7)}$$

which is strictly positive as $\Delta_{EU} > 0$. Thus, clearly $\gamma^* > 0.5$ and $\gamma^{\dagger*} > 0.5$. Further, γ^* can be equal to 1 or $\gamma^* < 1$, as it depends on the sign of $V'(\gamma = 1)$.

[15] To the best of our knowledge, Topkis (1978) was the first to prove this result.

To compare γ^* and $\gamma^{\dagger *}$ for intermediate cases, we compare the partial derivatives $\frac{\partial V(\gamma)}{\partial \gamma}$ and $\frac{\partial V(\gamma^\dagger)}{\partial \gamma^\dagger}$. We can implement a given mean-preserving spread via a sequence of pairwise mean-preserving spreads. It suffices to consider a single state k and pair (i, j) such that

$$\Delta_i(t) = \Delta_k + t_i, \qquad \Delta_j(t) = \Delta_k + t_j, \tag{A.1.8}$$

where $p_i t_i + p_j t_j = 0$, and $p_i + p_j = p_k$, leaving all states $\omega \neq k$ unchanged.

$$\frac{\partial V(\gamma)}{\partial \gamma} - \frac{\partial V(\gamma^\dagger)}{\partial \gamma^\dagger} = 2\lambda \Big[p_k R'\Big((2\gamma - 1)D_k\Big)D_k - p_i R'\Big((2\gamma - 1)D_i\Big)D_i - p_j R'\Big((2\gamma - 1)D_j\Big)D_j \Big]. \tag{A.1.9}$$

Thus, $\frac{\partial V(\gamma)}{\partial \gamma} - \frac{\partial V(\gamma^\dagger)}{\partial \gamma^\dagger} \geq 0$ if $zR'(\alpha z)$, with $\alpha > 0$ (which is guaranteed by $\gamma > 0.5$), is concave. $\frac{\partial^2 (zR'(\alpha z))}{\partial^2 z} = \alpha\Big(2R''(\alpha z) + \alpha z R'''(\alpha z)\Big)$, which is negative if $2R''(z) + zR'''(z) \leq 0$ for all $z \in \mathbb{R}$.

It is then easy to see that claim 4 holds. If $R''(z) = 0$ for $z < 0$, and $R''(z) \geq 0$ for $z > 0$, criterion 1 (regret-hedging) does not hold. Thus, we need $R''(z) < 0$ for $z < 0$. In addition, criterion 3 is achieved by imposing $R'''(z) \geq 0$ for $z < 0$, as well as $R''(z) = 0$ for $z > 0$, which implies $R'''(z) = 0$ for $z > 0$.

□

A.2 Proof of Proposition 2

Proposition 2 states that, holding the marginal distributions of lotteries A and B constant, a decrease in the concordance of the lotteries increases the rate at which the DM randomizes. Given that we impose regret-rejoice-risk monotonicity, a DM randomizes weakly more if the decrease in concordance induces a mean preserving spread in Δ_ω. In the following, we show that a decrease in the concordance induces a mean preserving spread in Δ_ω.

We prove the following lemma.

Lemma A.2.1 ("Concordance and Mean-Preserving Spreads of the Difference"). *Let* (X, Y) *and* $(X^\dagger, Y^\dagger)$ *be two bivariate random variables defined on the same finite state space* Ω *with common probabilities* $p_\omega > 0$*, and assume that the* marginal *distribution of* X *equals that of* $X^\dagger$*, and the marginal distribution of* Y *equals that of* $Y^\dagger$*.*

Define the state-wise differences as

$$\Delta_\omega := X_\omega - Y_\omega, \qquad \Delta_\omega^\dagger := X_\omega^\dagger - Y_\omega^\dagger, \quad \omega \in \Omega. \tag{A.2.1}$$

Denote F *and* $F^\dagger$ *the respective joint distributions. If* F *is more concordant than* $F^\dagger$ *in the concordance order, then* $\{\Delta_\omega^\dagger\}_{\omega\in\Omega}$ *is a mean-preserving spread of* $\{\Delta_\omega\}_{\omega\in\Omega}$*.*

Proof. Since we fix the marginals across joint distributions, by Theorem 3.8.2 in Müller and Stoyan (2002), F is more concordant than $F^\dagger$ if and only if $\mathbb{E}[\,g(X, Y)\,] \geq \mathbb{E}[\,g(X^\dagger, Y^\dagger)\,]$ for all supermodular functions $g : \mathbb{R}^2 \to \mathbb{R}$. Let $g(x, y) := h(x - y)$ for a twice differentiable $h(\cdot)$. Then $\frac{\partial^2 g(x,y)}{\partial x \partial y} = -\frac{\partial^2 h(x-y)}{\partial^2 (x-y)}$. Supermodularity of g is equivalent to $\frac{\partial^2 h(z)}{\partial^2 z} \leq 0$; i.e., $h(\cdot)$ is concave. Thus, the concordance order comparison implies that for every concave h, $\mathbb{E}[\,h(X - Y)\,] \geq \mathbb{E}[\,h(X^\dagger - Y^\dagger)\,]$. Since $\mathbb{E}[X - Y] = \mathbb{E}[X^\dagger - Y^\dagger]$, this coincides with the definition of $\{\Delta_\omega^\dagger\}_{\omega\in\Omega}$ being a mean-preserving spread of $\{\Delta_\omega\}_{\omega\in\Omega}$.

□

A.3 Feedback Effects

We consider the case in which the obtained outcome is not always perfectly informative about the realized state of the world, as this is the case in which the difference between the different feedback structures is relevant. For convenience, we restate the value function for the CR correlation structure (equation (2) in the main text) below.

$$V_{CR}(\gamma) = \gamma \sum_{\omega} p_\omega R\Big((2\gamma - 1)D^A_\omega\Big) + (1-\gamma) \sum_{\omega} p_\omega R\Big((2\gamma - 1)D^B_\omega\Big), \qquad \text{(A.3.1)}$$

where $D^A_\omega = \lambda\Big(u(x^A_\omega) - E\Big[u(x^B_\omega)|y = x^A_\omega\Big]\Big) + (1-\lambda)\Delta_{EU}$, and $D^B_\omega = \lambda\Big(E\Big[u(x^A_\omega)|y = x^B_\omega\Big] - u(x^B_\omega)\Big) + (1-\lambda)\Delta_{EU}$.

The first derivative is given by

$$\frac{\partial V_{CR}(\gamma)}{\partial \gamma} = \underbrace{2\gamma \sum_{\omega} p_\omega D^A_\omega R'\Big((2\gamma - 1)D^A_\omega\Big) + 2(1-\gamma) \sum_{\omega} p_\omega D^B_\omega R'\Big((2\gamma - 1)D^B_\omega\Big)}_{\text{(i) regret-hedging}} \qquad \text{(A.3.2)}$$

$$+ \underbrace{\sum_{\omega} p_\omega R\Big((2\gamma - 1)D^A_\omega\Big) - \sum_{\omega} p_\omega R\Big((2\gamma - 1)D^B_\omega\Big)}_{\text{(ii) asymmetric-feedback effect}}. \qquad \text{(A.3.3)}$$

As in AR, we can show that $\gamma^*_{CR} > (<)\ 0.5$ if $\Delta_{EU} > (<)\ 0$, by evaluating the derivative at $\gamma = 0.5$ and showing that it is positive if and only if $\Delta_{EU} > 0$.

For further discussion, it is convenient to reproduce the first derivative for the always-reveal case, where we add the subscript AR for clarity, and $D_\omega = [\lambda\Delta_\omega + (1-\lambda)\Delta_{EU}]$, which is independent of the DM's choice:

$$\frac{\partial V_{AR}(\gamma)}{\partial \gamma} = 2 \sum_{\omega \in \Omega} p_\omega D_\omega R'\Big((2\gamma - 1)D_\omega\Big). \qquad \text{(A.3.4)}$$

The first part of the derivative under CR, that we label "(i) regret-hedging", is similar to the derivative we obtain for AR. It captures the DM's regret-hedging problem, holding fixed the outcome information they receive.

However, there appears now also a second term labelled "(ii) asymmetric-feedback effect". This term captures that the DM's outcome information, and therefore their regret-rejoicing, now depends on the implemented lottery. If the implementation of one lottery shields the DM from regret-rejoice-risk, this creates an incentive for the DM to assign a higher probability to this option. An instructive example is a choice between a degenerate safe lottery and a risky lottery. If the safe option is implemented with certainty, the DM effectively obtains no outcome feedback, and hence the regret-rejoice-risk is eliminated. If the risky option is implemented, the DM receives outcome feedback and is subject to regret-rejoice-risk. This asymmetric-feedback effect induces the DM to assign a higher probability to the safe option.

We now compare the optimal mixtures under AR and CR. It is immediate that the distribution $\{(p_\omega, \Delta_\omega)\}_{\omega\in\Omega}$ is a mean preserving spread of $\gamma\{(p_\omega, \Delta_\omega^A)\}_{\omega\in\Omega} + (1-\gamma)\{(p_\omega, \Delta_\omega^B)\}_{\omega\in\Omega}$, for any $\lambda \in (0, 1]$. Formally, part (i) of $\frac{\partial V_{CR}(\gamma)}{\partial\gamma}$ is weakly greater than $\frac{\partial V_{AR}(\gamma)}{\partial\gamma}$ for $\gamma > 0.5$ if $\Delta_{EU} > 0$. Therefore, if it were not for the asymmetric-feedback effect, the DM would always randomize weakly less under CR than under AR. Intuitively, this captures the effect that partial outcome feedback dampens the sting of ex-post regret.

It follows that we can clearly identify the sign of the difference between $\frac{\partial V_{CR}(\gamma)}{\partial\gamma}$ and $\frac{\partial V_{AR}(\gamma)}{\partial\gamma}$ whenever part (ii) asymmetric-feedback effect is weakly positive. By regret-rejoice-risk aversion, this is the case whenever $\{(p_\omega, \Delta_\omega^A)\}_{\omega\in\Omega}$ weakly second order stochastically dominates $\{(p_\omega, \Delta_\omega^B)\}_{\omega\in\Omega}$. It is challenging to find conditions on the joint distribution of lotteries to ensure that this is indeed the case. When the two lotteries are independent and have equal expected utilities, the asymmetric-feedback effect favors the less risky of the two lotteries in the sense of a mean-preserving spread. However, this case is trivial in our model, as the DM chooses $\gamma = 0.5$ anyway. As soon as $\Delta_{EU} \neq 0$, the sign of the asymmetric-feedback effect generally depends on the shape of $u(\cdot)$. To see why, consider the specific case of the task we use in our experiment that we discuss in the next section. Reverting back to our illustrative example above, a special case is where one lottery is a degenerate safe option. In this case, the asymmetric-feedback effect always favors the safe option. The same holds true if one constructs compound lotteries from different choices between a degenerate and risky lottery.

A.4 Proof of Prediction 2

Denote $\Delta^{\mathcal{L}}(y = x_\omega^{\mathcal{L}})$ the counterfactual comparison when the DM chooses lottery $\mathcal{L}$ and observes its associated payoff y. In our experimental setup, there are four different counterfactuals:

- $\Delta^A(y = h + \pi) = u(h + \pi) - \frac{1}{2}[u(h) + u(\ell)]$
- $\Delta^A(y = \ell + \pi) = u(\ell + \pi) - \frac{1}{2}[u(h) + u(\ell)]$
- $\Delta^B(y = h) = \frac{1}{2}[u(h + \pi) + u(\ell + \pi)] - u(h)$
- $\Delta^B(y = \ell) = \frac{1}{2}[u(h + \pi) + u(\ell + \pi)] - u(\ell)$

Obviously, $\Delta^A(y = h + \pi) > \Delta^A(y = \ell + \pi)$ and $\Delta^B(y = h) < \Delta^B(y = \ell)$. Moreover, some simple algebra yields $\Delta^A(y = h + \pi) < \Delta^B(y = \ell)$, and $\Delta^A(y = \ell+\pi) > \Delta^B(y = h)$ if and only if $u(\ell+\pi) - u(\ell) > u(h+\pi) - u(h)$. Thus, if $u(\cdot)$ is concave, $\Delta^B(y = \ell) > \Delta^A(y = h+\pi) > \Delta^A(y = \ell+\pi) > \Delta^B(y = h)$. Given that $\Delta^A(y = h + \pi) + \Delta^A(y = \ell + \pi) = \Delta^B(y = h) + \Delta^B(y = \ell)$, the counterfactual distribution $\Delta^B(y = x_\omega^B)$ is a mean preserving spread of $\Delta^A(y = x_\omega^A)$. If $u(\cdot)$ is convex, the reverse is true. If $u(\cdot)$ is linear, the two counterfactual distributions are equivalent. Thus, if $u(\cdot)$ is concave or linear, we can clearly identify the sign of the asymmetric-feedback effect discussed above, and the DM randomizes weakly less under CR. $\square$

Appendix B: Agranov and Ortoleva (2017) and Agranov et al. (2023)

B.1 The Experimental Tasks

Table B.1.1 The Choice Tasks in Agranov and Ortoleva (2017)

	Lottery 1				Lottery 2			
Question	$\omega = 1$	$\omega = 2$	$\omega = 3$	$\omega = 4$	$\omega = 1$	$\omega = 2$	$\omega = 3$	$\omega = 4$
FOSD1	98	98	98	98	103	103	103	103
FOSD2	17	17	18	18	17	17	17	17
FOSD3	10	20	30	30	70	100	120	190
EASY1	23	23	30	30	5	5	5	31
EASY2	12	14	16	96	85	85	85	85
EASY3	100	100	100	100	20	20	20	101
HARD1	38	38	38	77	16	16	94	94
HARD2	10	10	90	90	32	45	45	56
HARD3	6	84	105	200	54	60	117	135
HARD4	13	30	51	81	19	32	38	86

Note. Each row presents one lottery pair. States of the world are equiprobable and implemented via the role of a fair four-sided die. Payoffs are in tokens, which are converted into US dollars at 20 tokens = $1.

In Agranov et al. (2023), the choices between dominant and dominated lotteries are such that the dominant lottery yields ($25, p;$5) with $p > 0.5$ and the dominated one yields ($25,$1 - p$ p;$5), with $p \in \{0.55, 0.60, 0.65, 0.70, 0.75, 0.80\}$. Uncertainty is implemented via a bingo cage with 20 balls numbered 1-20. Whenever the dominant lottery yields the high payoff, the dominated one yields nothing and vice versa.

Appendix C: Experimental Results

C.1 Main Experiment: Full Sample Analysis

Correlation Effects. Figure C.1.1 reproduces Figure 5 using the full sample. Black circles represent participants who pass the three sanity tests, while gray circles represent participants who do not. As can be seen, most choices of participants who are screened out are scattered around 50%, with a substantial amount below 50%. Averaging over all three correlation structures, screened out participants choose the dominant lottery with 53.9% probability. Considering the full sample, the average probability of choosing the dominant lottery is 70.7% under state-wise dominance, 66.6% under independence, and 63.6% under negative correlation. All pairwise comparisons are statistically significantly different at $p < 0.001$ (Wilcoxon signed-rank test). The fraction choosing the FOSD lottery with probability one is 22.7% under state-wise dominance, 10.6% under independence, and 8.7% under negative correlation. The two differences of interest remain statistically significant at $p < 0.001$ (two-sided test of proportions). Compared to the results presented in the main text, the correlation effects are smaller, consistent with attenuation bias.

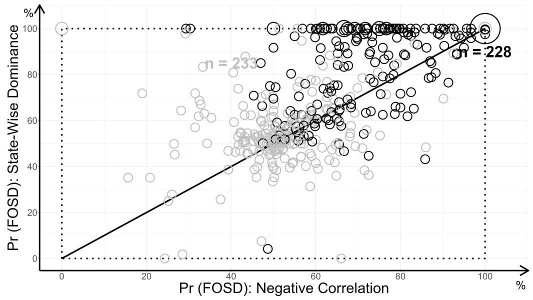

(a) State-Wise Dominance vs. Negative Correlation

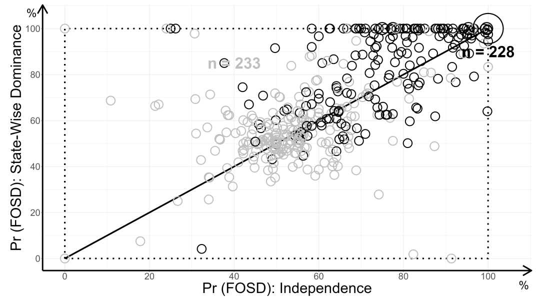

(b) State-Wise Dominance vs. Independence

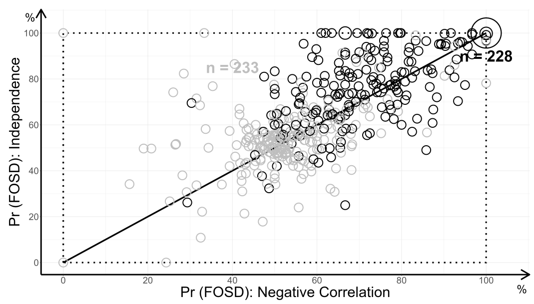

(c) Independence vs. Negative Correlation

Figure C.1.1 Full Sample: Randomization under Different Correlation Structures
The black circles are observations of participants who pass the sanity tests. The gray circles correspond to participants who do not pass these tests.

Figure C.1.2 displays the main results when we use the median instead of the mean mixture chosen by a given participant and correlation structure. Although there is now a higher prevalence of mixtures at focal probabilities such as 50% or 100%, overall patterns remain fairly similar. That is, a significantly larger share of mixtures falls above rather than below the 45-degree line. All three pairwise comparisons remain statistically significant at $p < 0.001$ (Wilcoxon signed-rank test).

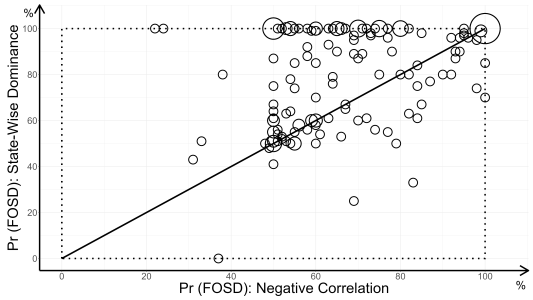

(a) State-Wise Dominance vs. Negative Correlation

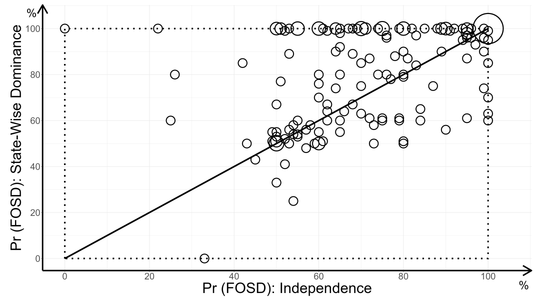

(b) State-Wise Dominance vs. Independence

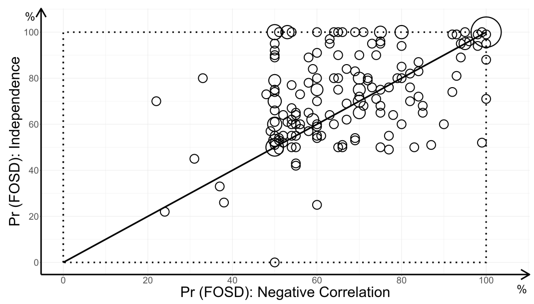

(c) Independence vs. Negative Correlation

Figure C.1.2 Full Sample – Median: Randomization under Different Correlation Structures

Lastly, Figure C.1.3 presents the mixtures for each lottery pair and individual under different correlation structures. The graphs appear noisier; however, substantially more observations still lie above than below the 45-degree line. Testing for a correlation effect at the lottery-pair level, we fail to reject the null in only one of the 18 comparisons ($p = 0.317$) and reject it at the $p = 0.070$ level in another. For the remaining 16 comparisons, we reject the null at $p \leq 0.020$.

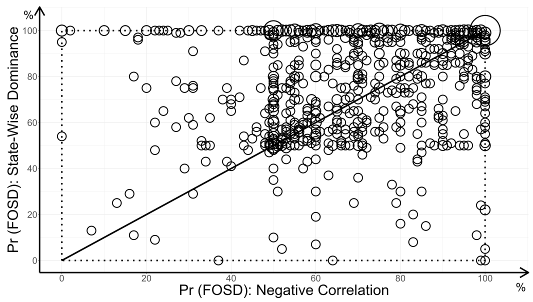

(a) State-Wise Dominance vs. Negative Correlation

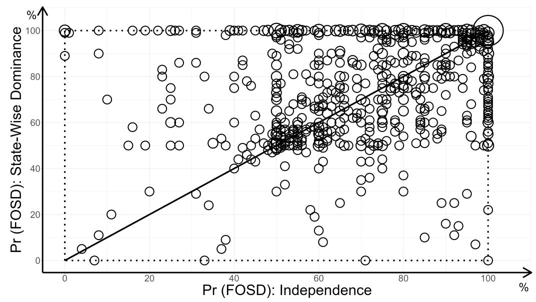

(b) State-Wise Dominance vs. Independence

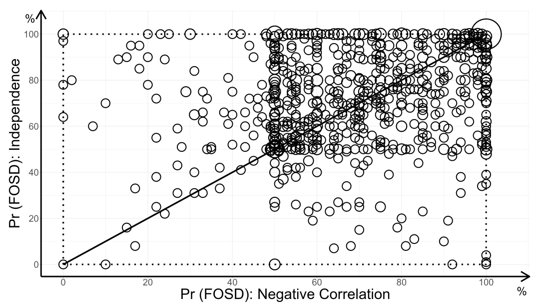

(c) Independence vs. Negative Correlation

Figure C.1.3 Full Sample – Disaggregated: Randomization under Different Correlation Structures

Ex-Post Regret. Considering the full sample, we find $H(110, 0) = 2.59 > H(110, 100) = 2.39 > H(10, 0) = 0.39 > H(10, 100) = -1.32$. Again, all pairwise comparisons are statistically significant at $p < 0.001$ (Wilcoxon signed-rank test). The differences are less pronounced as those reported in the main sample, consistent with attenuation bias.

Clustering Exercise. When conducting the clustering exercise using the entire sample, the obtained clusters are instable for three or more classes, meaning that different starting values produce different results. We do not face similar issues with the restricted sample, which suggest that many of the participants that we exclude with our sanity tests are indeed displaying erratic behavior, and, consequently, should be excluded from the main data analysis. The algorithm typically produces one class with a high number of participants ($\approx 85\%$) who fail the sanity tests and choose the dominant lottery with a probability close to 0.5

under all correlation structures. This class makes up close to 50% of the whole sample, which is similar to the overall fraction that does not pass the sanity tests (50.5%). Both the remaining two classes display correlation-sensitive PfR. The distinguishing feature is that, averaged across all correlation structures, one class chooses the dominant lottery with around 90% and the second class with around 65%.

When estimating four class models, we can qualitatively replicate the classes we obtain with our core sample. For instance, one estimation yields a class with 157 participants who choose the dominant lottery with 51.2% under state-wise dominance, with 46.8% under independence, and with 47.1% under negative correlation, respectively. Only 6.0% of the participants in this class pass the sanity tests. The remaining three classes are approximately similar to the three classes discussed in the main text. There is one class, with 106 participants displaying strong correlation sensitivity, one class with 76 participants mainly characterized by choice probabilities close to one, and another class with 122 participants displaying little correlation sensitivity and choosing the dominant lottery with approximately 60% probability. Again, the regret-hedging class displays the highest levels of ex-post regret and the second-highest average CRT score.

First Block. Participants were not informed about the correlation structures they would encounter later in the experiment. By considering the first block participants encountered, we can therefore alleviate concerns that could arise from the use of a within-subject design. In what comes next, we again focus on participants who pass the sanity tests. We find that that participants choose the FOSD lottery with 85.8% under state-wise dominance, with 77.7% under independence, and with 75.0% probability under negative correlation, respectively. These numbers are very similar to those presented in the main text. The comparison between state-wise dominance and negative correlation is statistically significant at $p < 0.001$, and the comparison between state-wise dominance and independence is statistically significant at $p = 0.008$. However, the comparison between independence and negative correlation now fails to reach statistical significance ($p = 0.303$, all p-values from Wilcoxon rank-sum tests), which is likely due to reduced statistical power.

C.2 Experiment 1: Selecting the Number of Clusters for the k-Means Exercise

In this section, we detail how we determine the optimal number of clusters k. Figure C.2.1 illustrates our diagnostic measures. We plot these diagnostics for 10,000 randomly selected starting values each, as k-means can be sensitive to the chosen starting values. The diagrams indeed suggest that, starting from $k > 6$, different starting values lead to different clusters.

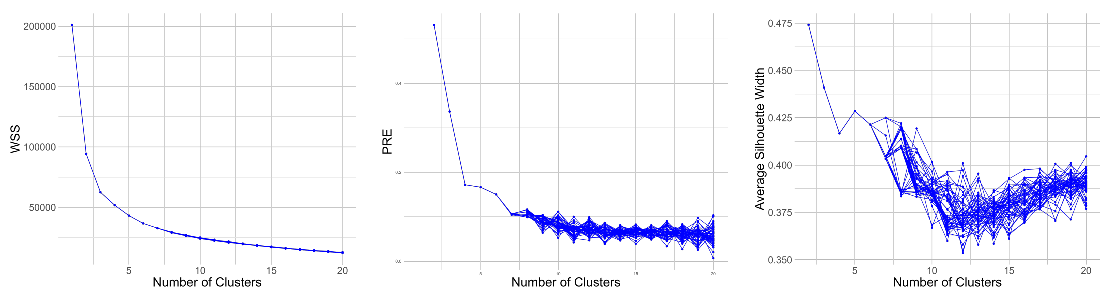

(a) Elbow Plot **(b)** PRE Plot **(c)** Silhouette Plot

Figure C.2.1 Criteria for Selecting the Number of Clusters

Figure C.2.1, Panel (a) plots the within-cluster sum of squares (WSS) against the number of clusters. The WSS decreases with the number of clusters, but higher numbers of clusters reduce parsimony. The elbow method (Thorndike, 1953) suggests visually inspecting the graph and selecting the number of classes based on where the "elbow" is located, i.e., where the marginal reduction in WSS levels off. In the diagram, the elbow is located at $k = 3$.

The proportional reduction in error (PRE) (Milligan and Cooper, 1985) measures the fraction by which the WSS drops when one more class is added to the estimation. It is given by $\text{PRE}(k) = 1 - \frac{\text{WSS}(k)}{\text{WSS}(k-1)}$ and is defined for $k \geq 2$. According to this method, the analyst selects the number of classes after which PRE stabilizes at a low level. The PRE, plotted in Figure C.4.1, Panel (b), begins to level off at $k = 4$.

The average silhouette width (Rousseeuw, 1987) can be used to measure how well separated different clusters are from each other. It is defined as follows. For each point i, denote $a(i)$ as its average distance to points in its own cluster and $b(i)$ as the average distance to points in the closest cluster (different from the cluster to which i is assigned). The silhouette for point i is defined as $s(i) = \frac{b(i)-a(i)}{\max\{a(i),b(i)\}} \in [-1, 1]$. The silhouette method selects k with the largest average silhouette $\left(\bar{s} = \frac{1}{n}\sum_i s(i)\right)$. As Figure C.4.1, Panel (b) shows, we achieve the greatest average silhouette width for $k = 2$, and the second-highest value for $k = 3$.

The different selection methods suggest two, three, or four classes. Overall, estimating $k = 3$ classes is well justified. This number of classes is also consistent with decision-theoretical reasons that suggest three classes.

C.3 Experiment 2: Display Effects

As in the main experiment, participants who do not pass the sanity tests choose the dominant lottery with a substantially lower probability, namely 54.3%, averaged over both correlation structures. When considering the entire sample, the average probability of choosing the FOSD lottery is 72.6% under state-wise dominance and 67.1% under negative correlation. Again, the difference remains statistically significant at $p < 0.001$ (Wilcoxon signed-rank test), although the size of the effect is somewhat attenuated.

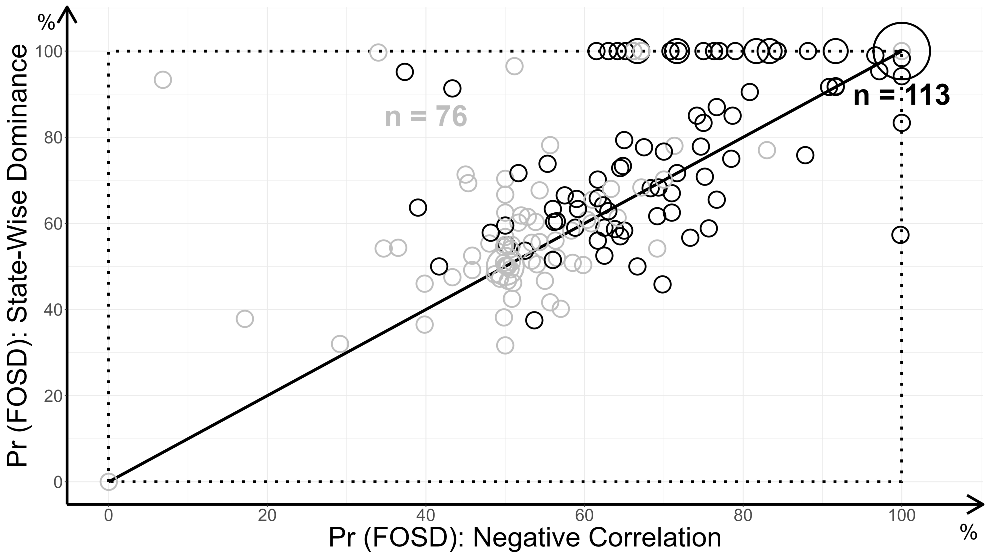

Figure C.3.1 Full Sample: Randomization under Different Correlation Structures
The black circles are observations of participants who pass the sanity tests. The gray circles correspond to participants who fail.

We can focus on the first block that participants encounter to test our main hypothesis between-subjects. Participants who pass the sanity tests choose the FOSD lottery with 86.6% probability under state-wise dominance and with 74.5% under negative correlation. The difference is high statistically significant at $p < 0.001$ (Wilcoxon rank-sum test).

C.4 Experiment 3: Feedback Effects

Similar to the first two experiments, participants who fail the sanity tests choose, on average, the FOSD lottery with 53.4% probability. Considering the whole sample, participants choose the dominant lottery with 73.1% under CR and with 69.4% under AR. The difference is statistically significant at $p = 0.003$ (Wilcoxon signed-rank test).

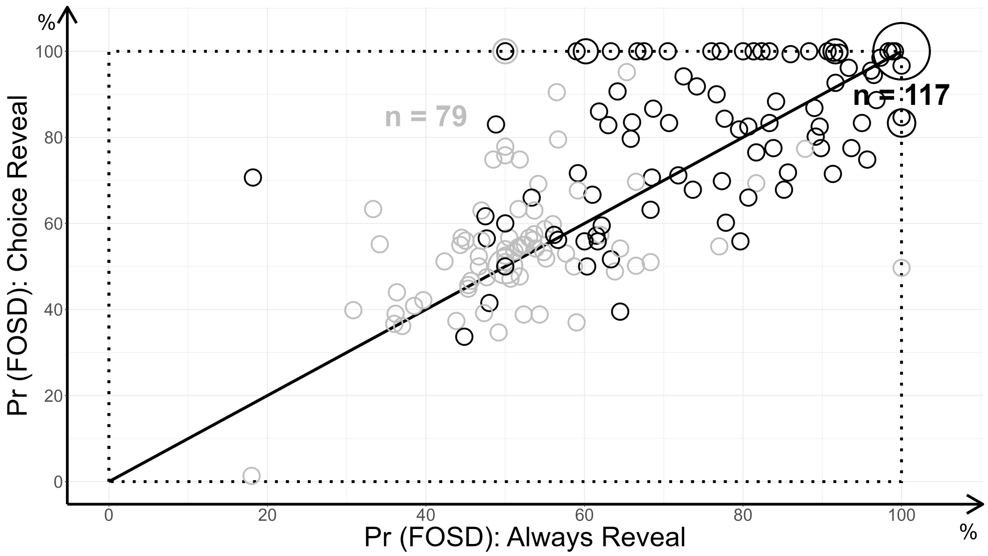

Figure C.4.1 Full Sample: Randomization under Different Feedback Structures
The black circles are observations stemming from participants who pass the sanity tests. The gray circles correspond to participants who fail.

Focusing on the first block only, participants passing the sanity tests choose the dominant lottery with 89.0% probability under CR and with 77.7% probability under AR. The difference is statistically significant at $p < 0.001$ (Wilcoxon rank-sum test).

Appendix D: Robustness of the Experimental Results

D.1 Econometric Robustness Exercises

Tables D.1.1, D.1.2, and D.1.3 report the results of several other robustness exercises. Using regression analysis, we control for round fixed effects, block fixed effects, order fixed effects, and individual fixed effects. As can be seen, the estimated coefficients change barely when controls are included. These findings indicate that our results remain robust, in both quantitative and qualitative terms, to ordering effects and potential experimental fatigue.

Table D.1.1 Additional Robustness Tests: Main Experiment

	(1) Baseline	(2) + Round FE	(3) + Block FE	(4) + Order FE	(5) + Individual FE	(6) All FE
	Dependent variable: `Weight assigned to the dominant lottery`					
Negative & Independence	-6.260***	-6.260***	-6.101***	-6.260***	-6.260***	-6.101***
	(0.995)	(0.996)	(1.000)	(0.996)	(1.024)	(1.030)
Negative	-4.472***	-4.472***	-4.476***	-4.472***	-4.472***	-4.476***
	(0.897)	(0.898)	(0.885)	(0.898)	(0.923)	(0.911)
Constant	85.871***	85.289***	84.947***	83.854***	86.164***	94.769***
	(1.182)	(1.279)	(1.351)	(2.374)	(0.663)	(1.087)
Round FE	No	Yes	No	No	No	Yes
Block FE	No	No	Yes	No	No	Yes
Order FE	No	No	No	Yes	No	Yes
Individual FE	No	No	No	No	Yes	Yes
Observations	4,104	4,104	4,104	4,104	4,104	4,104
Clusters	228	228	228	228	228	228
R-squared	0.038	0.039	0.039	0.050	0.463	0.465

Notes: OLS regressions with robust standard errors clustered at the individual level (in parentheses). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Round FE refer to the round within a given block in which a participant encountered a given choice task. Block FE control for the block (first, second, or third) in which a given participant makes a choice. Order FE control for the six possible orders in which participants could encounter the three correlation structures.

Table D.1.2 Additional Robustness Tests: Experiment 2 – Controlling for Display Effects

	(1) Baseline	(2) + Round FE	(3) + Block FE	(4) + Order FE	(5) + Individual FE	(6) All FE
	Dependent variable: `Weight assigned to the dominant lottery`					
Negative	-5.593***	-5.593***	-5.514***	-5.593***	-5.593***	-5.514***
	(1.385)	(1.388)	(1.389)	(1.386)	(1.446)	(1.453)
Constant	83.122***	84.084***	83.579***	79.567***	61.130***	104.214***
	(1.746)	(1.994)	(1.787)	(2.344)	(0.723)	(1.418)
Round FE	No	Yes	No	No	No	Yes
Block FE	No	No	Yes	No	No	Yes
Order FE	No	No	No	Yes	No	Yes
Individual FE	No	No	No	No	Yes	Yes
Observations	1,356	1,356	1,356	1,356	1,356	1,356
Clusters	113	113	113	113	113	113
R-squared	0.015	0.017	0.015	0.035	0.539	0.542

Notes: OLS regressions with robust standard errors clustered at the individual level (in parentheses). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Round FE refer to the round within a given block in which a participant encountered a given choice task. Block FE control for the block (first or second) in which a given participant makes a choice. Order FE control for the two possible orders in which participants could encounter the three correlation structures. Individual FE control for subject-specific factors by including an individual dummy for each participant.

Table D.1.3 Additional Robustness Tests: Experiment 3 – Feedback Effects

	(1) Baseline	(2) + Round FE	(3) + Block FE	(4) + Order FE	(5) + Individual FE	(6) All FE
	Dependent variable: `Weight assigned to the dominant lottery`					
AR	-4.043***	-4.043***	-4.011***	-4.043***	-4.043***	-4.011***
	(1.407)	(1.409)	(1.399)	(1.407)	(1.469)	(1.463)
Constant	85.315***	85.244***	85.566***	81.420***	91.188***	98.035***
	(1.598)	(1.781)	(1.775)	(2.379)	(0.734)	(1.576)
Round FE	No	Yes	No	No	No	Yes
Block FE	No	No	Yes	No	No	Yes
Order FE	No	No	No	Yes	No	Yes
Individual FE	No	No	No	No	Yes	Yes
Observations	1,404	1,404	1,404	1,404	1,404	1,404
Clusters	117	117	117	117	117	117
R-squared	0.008	0.008	0.008	0.034	0.500	0.501

Notes: OLS regressions with robust standard errors clustered at the individual level (in parentheses). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Round FE refer to the round within a given block in which a participant encountered a given choice task. Block FE control for the block (first or second) in which a given participant makes a choice. Order FE control for the two possible orders in which participants could encounter the three correlation structures. Individual FE control for subject-specific factors by including an individual dummy for each participant.