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

Johannes C. Hoelzemann *and* Moritz Loewenfeld

November 22, 2025

Abstract

A growing literature documents that many individuals deliberately delegate decisions to random devices such as coin flips, even when one option clearly dominates. We propose a novel explanation: randomization allows for hedging against risk of regret. In our model, a decision-maker engages in outcome-biased ex-post evaluations of their choices, projecting realized outcomes onto past decisions. Randomization ameliorates ex-post regret – if a bad outcome occurs, one can blame it on the coin flip. We conduct online experiments where participants choose mixtures over two lotteries, one of which first-order stochastically dominates. Holding marginal distributions fixed, we systematically vary the correlation of payoffs. As predicted by the model, participants randomize most under perfect negative correlation, less under independence, and least when one lottery dominates state-wise. In a clustering exercise, regret-hedgers emerge as the most prominent behavioral type. We further find that withholding outcome feedback on the non-chosen lottery decreases rates of randomization. Our model can rationalize puzzling findings in the literature, and we document the first direct evidence that randomization is a manifestation of deliberate regret-hedging.

JEL Classifications: C91, D81, D91.

Keywords: Preferences for randomization, regret, choice under risk, correlation sensitivity.

*

Johannes C. Hoelzemann, Universität Wien: https://www.johanneshoelzemann.com & jc.hoelzemann@gmail.com.
Moritz Loewenfeld, Universität Wien: https://sites.google.com/view/moritzloewenfeld & moritz.loewenfeld@univie.ac.at.

1 Introduction

A growing body of research documents that decision-makers (DMs) have a strong inclination toward deliberately stochastic choice. For instance, individuals often prefer to delegate a choice between lotteries to a coin toss, rather than making an explicit selection (Agranov and Ortoleva, 2017), which suggests deliberate *preferences for randomization (PfR)*. PfR are observed using a variety of elicitation methods (Agranov and Ortoleva, 2017; Feldman and Rehbeck, 2022; Agranov and Ortoleva, 2025), and appear to be a stable phenomenon across various domains (Agranov et al., 2023). PfR are not confined to the lab – they have been documented in the context of far-reaching decisions such as university applications, or whether to quit a job (Dwenger et al., 2018; Levitt, 2021). However, the underlying drivers of PfR remain unclear. The existing evidence violates not only expected utility theory, prominent behavioral models of risky choice (Tversky and Kahneman, 1992; Loomes and Sugden, 1982; Gul, 1991), but also the few models designed to accommodate PfR (Machina, 1985; Fudenberg et al., 2015; Cerreia-Vioglio et al., 2015).

In this paper, we propose and test a novel model in which PfR stem from a desire to hedge against risk of regret. We illustrate the main ideas with a simple thought experiment:

	1-3	4-6
A	110	10
B	0	100

Figure 1 Thought Experiment 1

Thought Experiment 1. Ann faces two options, A and B. A fair six-sided die is rolled. If the die shows 1-3, A pays \$110, while B pays nothing. In contrast, if the die comes up 4-6, A yields only \$10, whereas B pays \$100.

Will Ann choose A, B, or delegate the choice to a coin flip?

At first, choosing A seems preferable due to first-order stochastic dominance (FOSD). However, upon further consideration, Ann notices a dilemma. There is a 50% chance that the die shows 4-6. If this happens, Ann will regret choosing A – it would have been better to choose B. Choosing B is clearly worse than choosing A. On top of yielding lower payoffs, it also poses a 50% risk of regret due to the possibility of the die coming up 1-3. Fortunately, flipping a coin is a way out of Ann's quandary. No matter the outcome of the die roll, Ann can tell herself that she chose the right option with 50% probability. This mitigates Ann's risk of regret: if a bad outcome materializes, she can blame it on the coin flip.

In this paper, we formally develop the concept of "regret-hedging via randomization," derive novel predictions from comparative statics, test these through a series of carefully designed experiments, and offer a typology of randomizers.

The Model. The DM chooses a mixture in order to optimize their ex-post experience of regret and rejoice. At the heart of the model is an ex-post evaluation of the DM's choice. Intuitively, the DM judges their choice as if they should have known the realization of the random state of the world. For instance, if Ann chooses A and the die shows 4-6, she judges herself as if she should have known the outcome of the die roll, and thus feels regret because she made a "bad" choice. We model this as follows. At the ex-ante stage, the DM chooses a mixture, then observes the realized state, and engages in ex-post evaluation. In doing so, the DM projects the outcome information onto their decision at the ex-ante stage. The DM compares the expected utilities of both lotteries but inflates the weight given to the realized state. Depending on which lottery yields a higher value in this outcome-biased ex-post evaluation, the DM experiences regret or rejoice.

Crucially, our novel way of modeling ex-post regret makes the DM's objective function non-linear in probabilities. The choice they evaluate ex-post is a convex combination of both lotteries. Thus, randomizing allows the DM to decrease the magnitude of ex-post outcome comparisons and hedge against regret-rejoice risk. The DM may thus choose to randomize, even if one lottery is FOSD, as in thought experiment 1. If the DM does not project their outcome information onto their ex-ante decision, the model collapses to EUT. In this case, the DM evaluates their choice according to the unbiased expected utilities of lotteries A and B, and thus simply maximizes expected utility. Our model offers a psychologically rich perspective on how regret influences decision-making, and provides an intuitive rationalization of PfR.

The model also yields sharp comparative statics results. Holding constant the marginal distributions of lotteries A and B, the DM randomizes more when the correlation of payoffs becomes less concordant. To build intuition, consider the following thought experiment:

	1-3	4-6
A'	110	10
B'	100	0

Figure 2 Thought Experiment 2

Thought Experiment 2. Bob faces two options, A' and B'. A fair six-sided die is rolled. If the die shows 1-3, A' pays \$110, while B' pays \$100. In contrast, if the die comes up 4-6, A' yields only \$10, whereas B' pays nothing.

Will Bob choose A', B', or delegate the choice to a coin flip?

Clearly, Bob can now choose A' without risking any regret. No matter the outcome of the die roll, A' always fares better than B'. Note that, compared to Ann's thought experiment, the only difference is how payoffs are correlated across states. The marginal distributions of the lotteries remain constant.

This feature of the model – that rates of randomization are sensitive to the correlation of payoffs – rationalizes puzzlings findings in the literature. First, it can explain why Agranov and Ortoleva (2017) observe little randomization when one lottery dominates the other state-wise, while Agranov et al. (2023) observe frequent randomization when one lottery is FOSD but does not dominate state-wise. Second, it can rationalize why Agranov and Ortoleva (2017) observe more randomization in their "HARD" tasks than their "EASY" tasks, even when controlling for differences in expected utility.[1] As we discuss in more detail in Section 3 , the HARD tasks feature larger risks of regret than the EASY ones.

The prediction that the rate at which the DM randomizes depends on how payoffs are correlated also forms the basis for our main experiment. We present participants with pairs of lotteries, one of which is FOSD. They choose the probability with which the computer implements either lottery for them. Participants encounter each lottery pair in three correlation structures. In addition to perfect negative and perfect positive correlations (as in thought experiments 1 and 2), we include the intermediate case of independence, thus offering a more severe test of the model. The data confirm the main prediction of the model: participants indeed randomize most under perfect negative correlation, less under independence, and the least under state-wise dominance.

Our setup allows us to distinguish between three model classes: regret-hedging, perturbed expected utility (Machina, 1985; Fudenberg et al., 2015), and models that predict no randomization in our setting, which include models of rational mixing (Cerreia-Vioglio et al., 2015). To obtain a typology of randomizers, we conduct a clustering exercise using the k-means algorithm. Although the algorithm minimizes within-cluster distance and is thus non-parametric, we obtain three classes that approximate the three types predicted by theory surprisingly well. The regret-hedging class captures participants who randomize in a way that depends on the correlation of payoffs. With about 44% of participants assigned

[1] For details of the tasks, see the supplemental appendix.

to it, it emerges as the most frequent class. The remaining participants are split equally between the two other classes. The EUT class comprises participants who choose the dominant lottery with (near) certainty, while participants in the perturbed utility class randomize, but are unaffected by the correlation structure.

We present additional evidence that different rates of randomization across correlation structures are a manifestation of deliberate regret-hedging. We complement the choice data with a survey designed to construct measures of ex-post regret, and administer the cognitive reflection test (Frederick, 2005; Toplak et al., 2014). We find that participants who have been assigned to the regret-hedging class by the k-means algorithm score high in ex-post regret and cognitive reflection. This constitutes correlational evidence supporting deliberate regret-hedging. In addition, we conduct a second experiment in which we rule out that our main results are driven by potential confounds related to the choice display.

We conduct one further mechanism experiment in which we test feedback effects. If randomization is used to hedge against ex-post regret, the content of outcome feedback influences rates of randomization. We consider two feedback structures: DMs obtain either feedback on both the obtained and forgone outcome, or on the obtained outcome only. We consider a setting as close as possible to the previous experiments. However, we focus on independence, as the obtained payoff in perfectly informative about the forgone payoff under perfect positive or negative correlation. As our framework explicitly models regret as stemming from information projection, it can easily accommodate different feedback structures. Withholding feedback on the forgone outcome dampens ex-post regret-rejoice, which reduces the need to hedge against regret-rejoice risk by randomizing. In the experiment, we indeed find that participants randomize less under partial outcome feedback.

2 Related Literature

Our paper relates to and connects several strands of literature, which we discuss in turn.

Stochasticity as Noise. A number of studies, starting with Mosteller and Nogee (1951), have documented stochasticity in choice (Mosteller and Nogee, 1951; Tversky, 1969; Camerer, 1989; Hey and Orme, 1994; Hey, 2001). In these studies, participants are confronted with the same choice task more than once, usually separated by a number of other tasks. Participants often make different choices when confronted with the same choice task, which is often interpreted as resulting from stochastic preferences or some form of decision noise, for instance due to bounded rationality (Loomes and Sugden, 1995; Manzini and Mariotti, 2014).

Experiments on PfR. Several studies investigate whether stochastic choice can be a *deliberate* expression of preferences. Agranov and Ortoleva (2022) provides a review of this nascent literature. A few earlier studies have documented stochastic choices when participants choose mixtures explicitly (Loomes, 1998; Sopher and Narramore, 2000; Rubinstein, 2002). More recently, Agranov and Ortoleva (2017) document deliberate randomization and ascribe it to preferences. They ask participants to choose repeatedly, explicitly alerting them that they make repeated choices between the same two lotteries. They also find that several participants are willing to incur a small cost to flip a coin, rather than choosing one lottery by themselves. Further, as discussed above, rates of randomization differ between different types of tasks. Cettolin and Riedl (2019) distinguish deliberate preferences for randomization from stochasticity due to incomplete preferences in choices between risky and ambiguous lotteries. Agranov et al. (2023) find that randomization is a stable trait across different domains of decision-making. Feldman and Rehbeck (2022) elicit participants' preferred mixture over binary lotteries, and find that these mixtures correlate with randomization decisions measured by repeated choices. Importantly, they rule out that mixing behavior is driven by experimenter demand or noise. Agranov and Ortoleva (2025) elicit PfR in a choice-list format and find that participants display PfR for a wide range of parameter values. Chew et al. (2022) provide evidence that suggests that multiple switching in multiple-choice lists are caused by PfR. Toussaert (2024) studies PfR in the context of choices between holiday locations, and rationalizes randomization as reflecting a demand for surprise. PfR have also been documented for far-reaching choices outside the laboratory. Dwenger et al. (2018) study university applications in Germany, and find that the rankings of universities submitted by about roughly a third of applicants are consistent with PfR. Follow-up questionnaires provide direct evidence that many applicants randomize deliberately. Levitt (2021) finds that many people follow through with the suggestion of the coin flip, even for decisions such as quitting a job, ending a relationship, or starting a business.

Models Accommodating PfR. Several existing models can accommodate some forms of PfR. First, though not conceived to explain PfR, a number of models that are non-linear in probabilities allow for some mixing behavior (Kahneman and Tverski, 1979; Quiggin, 1982; Chew et al., 1991; Tversky and Kahneman, 1992). Cerreia-Vioglio et al. (2019) provide a representation of deliberate stochastic choice based on a notion of rational mixing. Individuals have preferences defined over the marginal distribution of payoffs that results from their mixture. Importantly, these preferences always satisfy FOSD. Models satisfying rational mixing are cautious expected utility (Cerreia-Vioglio et al., 2015), as well as rank-dependent models (Quiggin, 1982; Tversky and Kahneman, 1992). A distinct class of preferences are perturbed utility models (Machina, 1985; Fudenberg et al., 2015). In this class, DMs can be modeled as having a convex cost of choosing mixture probabilities. This cost may capture the DM's uncertainty about their preferences, inattention or cost of implementing their preferences (Fudenberg et al., 2015). The closest theoretical paper to ours is Heydari (2024). In his model, the DM might randomize to reduce responsibility for outcomes and thereby regret. Our paper differs in important ways. First, in Heydari (2024), randomization results from an exogenous aversion to responsibility, whereas it ultimately results from ex-post projection in our model. Second, comparative statics discussed in Heydari (2024) are confined to differences in responsibility aversion, whereas we focus on comparative statics that pertain to the DM's choice environment, i.e., the correlation of payoffs and the feedback structure.

Existing Models: Empirical Evidence. In their seminal paper, Agranov et al. (2023) consider different models of mixing in light of the existing evidence, and conclude that none of them can organize the evidence in a satisfactory way. Models that are nonlinear in probabilities would require a probability weighting function that is quite specific and inconsistent with the common inverse-S shape. Models of rational mixing satisfy FOSD and therefore cannot rationalize the high rates of randomization when one lottery is dominant. While perturbed utility models can organize their data, Agranov et al. (2023) argue that these models are unsatisfactory because they provide no rationale for mixing but "simply posit that mixing is chosen because the subject has a utility for mixing"(Agranov et al., 2023, p.2573). To this, we would add that perturbed utility models make no difference between state-wise dominance and mere FOSD. Perturbed utility models thus cannot rationalize why Agranov and Ortoleva (2017) report such high rates for non-state-wise FOSD, whereas Agranov and Ortoleva (2017) find hardly any randomization for state-wise dominance. As we discuss in more detail in Section 3.4 regret-hedging rationalizes the existing evidence and provides an intuitive rationale for why people randomize.

Connections to Regret Theory. Our model builds on insights from and contributes to a large theoretical and experimental literature on regret theory (Loomes and Sugden, 1982; Bell, 1982, 1983; Quiggin, 1990, 1994; Sarver, 2008; Diecidue and Somasundaram, 2017; Gollier, 2020; Loewenfeld and Zheng, 2024) and correlation sensitivity more generally (Fishburn, 1982; Lanzani, 2022; Loewenfeld and Zheng, 2025; Chen et al., 2025). Our main contribution to this literature is that we offer a novel way of modeling regret. Whereas the psychology literature typically describes regret as an emotion that arises when "a *decision* [emphasis added] appears to be wrong in retrospect" (Zeelenberg et al., 2000, p.521), regret is usually modeled as resulting from a counterfactual comparison between the realized and forgone *outcome* (Loomes and Sugden, 1982; Bell, 1982, 1983; Sarver, 2008; Gollier, 2020). The importance of the DM's choice is often acknowledged, for instance, in the concept of "choiceless utility" used in Loomes and Sugden (1982). However, this approach raises a question: If the DM is fully aware of the distribution of outcomes when they make a choice, why would they regret this choice when it yields a bad outcome? To the best of our knowledge, we are the first to model regret explicitly as resulting from an ex-post assessment of the DM's choice. In our model, the DM regrets their choice because they evaluate it in hindsight. This approach is consistent with a fast-growing literature showing that individuals often judge others' choices with the benefit of hindsight (Baron and Hershey, 1988; Fischhoff, 1975; Gurdal et al., 2013; Danz et al., 2015; Loewenfeld, 2025). In particular, we build on Loewenfeld (2025) who models outcome biased ex-post evaluations in a setting of delegated risk taking. We believe that this approach to modeling regret yields two main advantages, in addition to providing psychological depth. First, it naturally produces preferences for randomization without the need for explicit assumptions. Second, the model seamlessly accommodates different feedback structures, whereas existing approaches typically need to impose strong additional assumptions (Bell, 1983; Gabillon, 2020; Zheng, 2021).

3 A Model of Regret-Hedging

3.1 The Setup

We adopt the description of risk as defined by states of nature following Savage (1954) that is common in regret and salience theory (Loomes and Sugden 1982, 1987, Bordalo et al. 2012). There is a finite set of states Ω, and a finite set of outcomes $X \in \mathbb{R}$. Each state ω occurs with an objective probability p_ω. Slightly abusing terminology, a lottery $\mathcal{L}$ is a function that assigns a real-valued payoff $x_\omega^{\mathcal{L}} \in X$ to each possible state of nature. In what follows, we consider binary choice sets, with $\mathcal{L} \in \{A, B\}$. The DM chooses the probability γ with which a random device implements lottery A, and lottery B is implemented with probability $1-\gamma$. In turn, the DM receives the realized payoff of the lottery that is implemented by the random device.

For expositional purposes, it is useful to distinguish the *ex-ante* and the *ex-post* stages. At the ex-ante stage, the DM chooses γ, and a lottery is implemented. At the ex-post stage, the DM learns the realized state of the world, evaluates the decision taken at the ex-ante stage, and experiences the consequence of this decision. If they conclude that they made a bad choice, they experience regret, an aversive emotion. If the DM concludes that they made a good choice, they experience rejoice. At the outset, the DM chooses γ taking the expected regret and rejoice into account.

We assume that the DM's ex-post evaluation of their ex-ante choice is outcome-biased. Following Madarász (2012) and Loewenfeld (2025), we assume that the DM (partially) projects the outcome information onto the ex-ante stage. This induces the DM at the ex-post stage to evaluate their choice as if they could have anticipated the realized state of the world. More formally, suppose that state ω materializes. The DM then computes the expected utility of lottery $\mathcal{L}$ but does so in an outcome-biased way. We assume that they inflate the decision weight of the realized state of the world to $p_\omega + \lambda(1-p_\omega)$, where $\lambda \in [0,1]$ is the DM's degree of OB. At the same time, they deflate the weights of the non-realized states $\omega' \neq \omega$ uniformly, to $(1-\lambda)p_{\omega'}$, which ensures that total weights sum to one. This formulation of OB is equivalent to the way information projection is modeled in Madarász (2012). It also captures the most prominent mechanism proposed in the psychology literature: in the seminal paper on the topic, Baron and Hershey (1988) suggest that OB arises because the realized outcome draws disproportionate attention.

The outcome-biased expected utility of lottery $\mathcal{L}$ in state ω, $\widetilde{EU}_\omega^{\mathcal{L}}$, can be conveniently written as:

$$\widetilde{EU}_\omega^{\mathcal{L}} = \lambda \underbrace{u(x_\omega^{\mathcal{L}})}_{\text{Realized outcome}} + (1-\lambda) \underbrace{\sum_{\omega \in \Omega} p_\omega u(x_\omega^{\mathcal{L}})}_{\text{Ex-ante expected utility}}$$

That is, the outcome-biased expected utility is a convex combination of the realized utility and the ex-ante expected utility of lottery $\mathcal{L}$, with weights determined by the DM's degree of OB λ. When $\lambda = 0$, the expression reduces to standard expected utility. When $\lambda = 1$, the DM assigns full weight to the realized outcome. The outcome-biased expected utility of the mixture γ is given by:

$$\widetilde{EU}_\omega^{\gamma} = \lambda\Big[\gamma u(x_\omega^A) + (1-\gamma)u(x_\omega^B)\Big] + (1-\lambda)\Big[\gamma EU^A + (1-\gamma)EU^B\Big],$$

where $EU^{\mathcal{L}} = \sum_{\omega \in \Omega} p_\omega u(x_\omega^{\mathcal{L}})$.

To assess whether they could have done better, the DM compares their chosen mixture γ to the counterfactual mixture $1-\gamma$. After choosing the mixture γ, the DM experiences regret or rejoice in state ω:

$$R\Big((2\gamma-1)[\lambda\Delta_\omega + (1-\lambda)\Delta_{EU}]\Big),$$

where $R : \mathbb{R} \to \mathbb{R}$ is an increasing, continuous, thrice differentiable function, $\Delta_\omega = u(x_\omega^A) - u(x_\omega^B)$, and $\Delta_{EU} = EU^A - EU^B$.

Finally, the DM chooses γ to maximize the expected regret-rejoice:

$$\gamma = \arg\max_{\gamma\in[0,1]} \sum_{\omega\in\Omega} p_\omega R\Big((2\gamma-1)[\lambda\Delta_\omega + (1-\lambda)\Delta_{EU}]\Big).$$

A couple of remarks are in order: First, the model nests EUT when $\lambda = 0$. In this case, the DM evaluates their choice based on the unbiased expected utility of lotteries A and B, regardless of the realized state. They therefore choose the lottery with the higher expected utility with probability 1. In this case, the DM might only randomize in the knife-edge case where lotteries A and B have the same expected utility. Second, the choice of the counterfactual, $1-\gamma$, merits a brief discussion. It is straightforward that this is the unique formulation that satisfies the following requirements. The counterfactual is a linear function of γ. Moreover, whenever the DM chooses one lottery with probability 1, the DM compares (100% of) the chosen lottery to (100% of) the non-chosen lottery.[2] We therefore argue that this counterfactual is the natural choice in our setting.

Finally, when γ is restricted to 0 or 1, one obtains an expression reminiscent of original regret theory (Loomes and Sugden 1982). In our model, the DM prefers lottery A whenever

$$\sum_{\omega\in\Omega} p_\omega \Big[R\Big(\lambda\Delta_\omega + (1-\lambda)\Delta_{EU}\Big) - R\Big(-\lambda\Delta_\omega - (1-\lambda)\Delta_{EU}\Big)\Big] \geq 0.$$

In original regret theory, adjusted to our notation, the DM prefers lottery A whenever

$$\Delta_{EU} + \sum_{\omega\in\Omega} p_\omega \Big[R\Big(\Delta_\omega\Big) - R\Big(-\Delta_\omega\Big)\Big] \geq 0.$$

These models differ in a few important aspects. First, in original regret theory, regret arises from an ex-post comparison of *outcomes*, and is linear in probabil-

[2] If the counterfactual is a linear function of γ, it can be written as $c+b\gamma$, where c is a constant. Moreover, the second requirement implies that if $\gamma = 0$, $\gamma - c - b\gamma = -1$, or $c = 1$. When $\gamma = 1$, the condition requires $\gamma - 1 - b\gamma = 1$, or $b = -1$. This implies the counterfactual $1-\gamma$.

ities. In our model, regret arises from an outcome-biased evaluation of *choices*. It is this feature of the model that makes regret-rejoice nonlinear in probabilities. Second, in both models, the DM trades off the ex-ante expected utility with ex-post regret. In original regret theory, this tradeoff is additively separable. In our model, however, the difference in ex-ante expected utility enters the regret function directly, and is therefore not additively separable from outcome regret (unless $R(\cdot)$ is linear). Finally, in our model, the extent to which anticipated regret shapes preferences is governed by the outcome-bias parameter λ.

3.2 Regret-Hedging and Aversion to Risk of Regret-Rejoice

To develop our notion of regret-rejoice-risk aversion, we impose the following intuitive criteria on randomization behavior:

Definition 1 ("Regret-Rejoice-Risk Aversion")**.** *We say that a DM satisfies regret-rejoice-risk aversion, if they display the following behavior:*

1. *Regret-hedging. Whenever* $\Delta_{EU} = 0$ *and* $\Delta_\omega > 0$ *for some states but* $\Delta_\omega < 0$ *for some other states, the DM strictly prefers* $\gamma = 0.5$*, for all* $\lambda \in (0, 1]$.

2. *Payoff monotonicity. If the payoff of lottery A in some state k is changed by* π*,* $\gamma(\pi)$ *is non-decreasing in* π*.*

3. *Regret-rejoice-risk monotonicity. The DM randomizes weakly more for a mean-preserving spread in* Δ_ω*.*

Regret-hedging states that a DM who chooses between two lotteries with the same expected utility maximizes expected regret-rejoice by fully randomizing. Intuitively, a regret-hedger with a utility function as defined above generally trades off two motives: a desire to minimize risk in regret-rejoice by randomizing and a desire to choose the lottery with the higher expected utility. When both lotteries yield the same expected utility, the second motive is muted and the

DM fully randomizes. *Payoff monotonicity* is a minimum regularity requirement. Whenever a lottery improves unambiguously, the DM assigns a higher probability to that lottery. Finally, *regret-rejoice-risk monotonicity* imposes regularity on regret-hedging. A mean-preserving spread in Δ_ω implies an increase in regret-rejoice-risk while the difference in expected utilities remains constant. If the DM randomizes in order to hedge against regret-rejoice-risk, they randomize more if the regret-rejoice-risk increases.

We show below that the following assumptions on the function $R(\cdot)$ of a regret-rejoice-risk averse DM are sufficient to generate these three criteria on randomization behavior.

Proposition 1 ("Functional Form Restrictions"). *The following assumptions on $R(\cdot)$ generate regret-rejoice-risk aversion:*

1. *Regret-hedging implies $\mathbb{E}[R(Y)] < R\big(\mathbb{E}[Y]\big)$ for every random variable Y with $\mathbb{E}[Y] = 0$, $\Pr(Y > 0)\Pr(Y < 0) > 0$.*
2. *Payoff monotonicity implies that $\frac{\partial R(z)}{\partial z} + z\frac{\partial^2 R(z)}{\partial z} \geq 0$, $\forall z \in \mathbb{R}$.*
3. *Regret-rejoice-risk monotonicity implies that $2\frac{\partial^2 R(z)}{\partial z} + z\frac{\partial^3 R(z)}{\partial z} \leq 0$, $\forall z \in \mathbb{R}$.*
4. *Sufficient conditions for regret-hedging, payoff monotonicity, and regret-rejoice-risk monotonicity: $\frac{\partial^2 R(z)}{\partial z} < 0$ and $\frac{\partial^3 R(z)}{\partial z} \geq 0$ for $z < 0$ (concavity in the regret-domain), and $\frac{\partial^2 R(z)}{\partial z} = 0$ for $z > 0$ (linearity in the gain domain).*

We relegate the proof to the supplemental appendix but briefly discuss here how we proceed. Claim 1 is straightforward. To prove claim 2, we use monotone comparative statics (Topkis, 1978; Milgrom and Shannon, 1994). The condition $\frac{\partial R(z)}{\partial z} + z\frac{\partial^2 R(z)}{\partial z} \geq 0$ guarantees that the value function $V(\gamma, \pi)$ is super-modular. To prove claim 3, we first show that the DM chooses $\gamma > 0.5$ whenever $\Delta_{EU} > 0$. We then consider two lottery pairs, (A, B) and $(A^\dagger, B^\dagger)$. We impose that the random variable $F^\dagger = \{(p^\dagger_\omega, \Delta^\dagger_\omega)\}_{s=1}^{S}$ is a mean preserving spread of $F =$

$\{(p_\omega, \Delta_\omega)\}_{s=1}^S$. This ensures that the difference in expected utilities is the same, but the risk of regret-rejoice is higher for the lottery pair $(A^\dagger, B^\dagger)$. We then compare the marginal gain in expected regret-rejoice from increasing $\gamma \in (0.5, 1]$ across the two lottery pairs, and show that the marginal gain is smaller for the lottery pair lottery pair $(A^\dagger, B^\dagger)$ for $\lambda \in (0, 1]$, if $2\frac{\partial^2 R(z)}{\partial z} + z\frac{\partial^3 R(z)}{\partial z} \leq 0$, for all $z \in \mathbb{R}$.

Given that claims 1-3 hold, it is easy to see that claim 4 holds as well. If $\frac{\partial^2 R(z)}{\partial z} = 0$ for $z < 0$, and $\frac{\partial^2 R(z)}{\partial z} \geq 0$ for $z > 0$, *regret-hedging* does not hold. Thus, we need $\frac{\partial^2 R(z)}{\partial z} < 0$ for $z < 0$. In addition, *regret-rejoice-risk monotonicity* is achieved by imposing $\frac{\partial^3 R(z)}{\partial z} \geq 0$ for $z < 0$, as well as $\frac{\partial^2 R(z)}{\partial z} = 0$ for $z > 0$, which implies $\frac{\partial^3 R(z)}{\partial z} = 0$ for $z > 0$, giving the desired property.

3.3 Randomization and the Correlation of Payoffs

We now discuss which aspects of the choice situation lead to *relatively more or less* randomization. Proposition 2 highlights the key feature of regret-hedging – the extent to which a DM wishes to randomize depends on the *correlation* of payoffs across states.

Proposition 2 (“More Randomization Under Less Concordant Correlation”). *Consider two lottery pairs (A, B) and $(A^\dagger, B^\dagger)$ such that A and $A^\dagger$, as well as B and $B^\dagger$, share the same marginal distribution. Denote γ^* and $\gamma^{\dagger *}$ the corresponding optimal mixtures chosen by a regret-hedging DM with degree of OB λ. If the payoffs of (A, B) are more concordant than those of $(A^\dagger, B^\dagger)$, then $\gamma^* \geq \gamma^{\dagger *} > 0.5$, for $\lambda \in (0, 1]$.*

We relegate the proof of Proposition 2 to the supplemental appendix. We show that a decrease in concordance implies a mean-preserving spread in Δ_ω. The result then follows immediately from regret-rejoice-risk monotonicity.

The proposition states that, holding the marginal distribution of the lotteries A and B constant, a decrease in the concordance of the lotteries' payoffs increases the rate at which the DM randomizes. Intuitively, the less concordant two lotteries are, the higher the propensity for ex-post regret – and thus the stronger the incentive for regret-hedging.[3]

3.4 Correlation-Sensitivity Rationalizes Puzzling Findings

Our model can rationalize several findings in the literature. First, it rationalizes why Agranov and Ortoleva (2017) found that participants randomize more in their HARD tasks than their EASY tasks, also when controlling for differences in expected utility.[4] A closer look at these tasks reveals that HARD tasks involve a much higher risk of regret than EASY tasks. When a participant chooses the expected value maximizing lottery, the maximum amount of experimental currency they can forgo is 1 token for tasks EASY 1 and 3, and 11 tokens for EASY 2. This compares to 22 tokens for HARD 1, 35 for HARD 2, and 65 tokens for HARD 3. HARD 4 features two lotteries with the same expected value. Our model of regret-hedging predicts higher rates of randomization for the HARD than the EASY tasks.

Moreover, Agranov and Ortoleva (2017) find that participants randomize little when one lottery dominates the other state-wise, whereas Agranov et al. (2023) find relatively high rates of randomization when one lottery is FOSD but does not yield a higher payoff in every state. Clearly, a regret-hedger would never randomize when one lottery is state-wise dominant. Choosing the dominant lottery with probability one eliminates all risk of regret. Since $R(\cdot)$ is increasing, randomizing can only decrease the DM's expected regret-rejoice. However, whenever the dominated lottery yields a higher payoff in at least one state, a regret-hedger

[3] Concordance is a measure of the dependence between two variables. When marginal distributions are held fixed, higher concordance implies higher covariance, but the reverse is not generally true. Well-known measures of concordance include Kendall's τ and Spearman's ρ, both of which are rank-based measures that order joint distributions from perfectly discordant ($\tau = \rho = -1$) to perfectly concordant ($\tau = \rho = 1$).

[4] See the supplemental appendix for a detailed description of these choice problems.

with $\lambda > 0$ may randomize, which can explain the higher rates of randomization observed by Agranov et al. (2023).

3.5 Testing Correlation-Sensitivity in Preferences for Randomization

While the patterns documented in Agranov and Ortoleva (2017) and Agranov et al. (2023) are consistent with preferences for randomization depending on the correlation structure, their experiments were not designed to isolate this property. In this paper, we design an experiment to identify correlation-sensitivity in PfR, while carefully controlling for potential confounds. Participants are presented with lottery pairs and are asked to choose the probability with which the computer implements either lottery. Each lottery pair includes one FOSD lottery. Participants encounter choice tasks that hold constant the marginal distributions of lotteries and differ only in how payoffs are correlated across states. In particular, we consider three correlation structures: state-wise dominance, independence, and perfect negative correlation.

Our main predictions follow immediately from Proposition 2. Prediction 1 states that we expect participants to display the lowest rates of randomization under state-wise dominance, more randomization under independence, and the highest rates under perfect negative correlation.

Prediction 1 ("Rates of Randomization Depend on the Correlation Structure")**.** *Denote γ participant's probability assigned to the FOSD lottery:* $\gamma(\textit{State-Wise Dominance}) > \gamma(\textit{Independence}) > \gamma(\textit{Perfect Negative Correlation})$.

We employ this setting for a number of reasons. First, an advantage of using FOSD is that lower mixture probabilities assigned to the dominant lottery are synonymous with more randomization. Models that are linear in probability such as EUT predict that participants choose the dominant lottery with probability one under *any* correlation structure, as do models satisfying FOSD, such as prospect theory (Tversky and Kahneman, 1992). In our setting, this holds true even for correlation-sensitive models that do not generally obey FOSD, e.g., regret or salience theory (Loomes and Sugden 1982, 1987, Bordalo et al. 2012), because the dominant lottery is always regret dominant (Quiggin, 1990).[5]. In the regret-hedging model, an outside observer knows that $\Delta_{EU} > 0$, regardless of the curvature of $u(\cdot)$. The regret-hedging model thus predicts $\gamma \in (0.5, 1]$. Perturbed utility models (Fudenberg et al., 2015), likewise predict that $\gamma \in (0.5, 1]$. Thus, our setting allows for a clear interpretation of mixtures chosen by participants.

Furthermore, the setting allows for a clear separation of regret-hedging from other random choice models. Figure 3 illustrates the predictions of regret-hedging and those of other prominent models graphically. We plot the probability with which the FOSD lottery is chosen under the perfect negative correlation on the x-axis against the probability with which it is chosen under both state-wise dominance and independence on the y-axis. The predictions of regret-hedging are illustrated in red color. The extreme case of state-wise dominance is illustrated by the dark red line, while the red-shaded area below corresponds to the intermediate case of independence. While choice probabilities are always predicted to be above the 45-degree line, for state-wise dominance, they equal one. Models of deliberate stochastic choice (Cerreia-Vioglio et al., 2015), illustrated in green color, obey FOSD and thus predict no randomization in the present setting. Finally, revealed perturbed utility models (Machina, 1985; Fudenberg et al., 2015), highlighted in blue color, predict that choice probabilities are greater than 0.5 and lie on the 45-degree line.

[5] With equiprobable states, lottery A regret dominates lottery B if, for every state k in which $x_k^B > x_k^A$, there is a state j such that $x_j^A \geq x_k^B$ and $x_j^B \leq x_k^A$, with at least one strict inequality for one state.

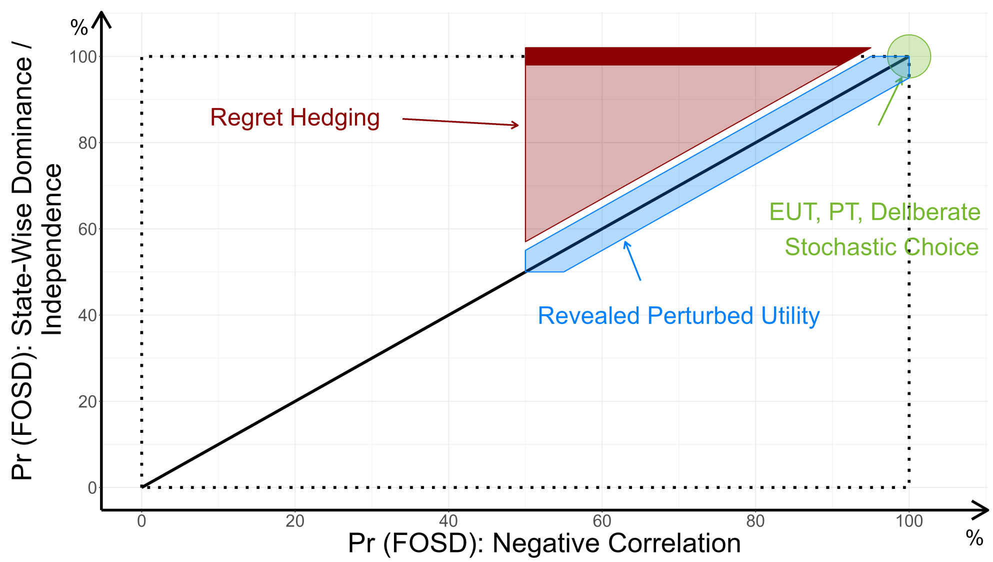

Figure 3 Predictions of Different Theories

4 Main Experiment: Correlation-Sensitivity in Preferences for Randomization

4.1 Details and Procedures

In the following, we lay out the details of the experiment.

Lottery Parameters. The two lotteries in each choice task are structured as follows. Both lotteries yield a high payoff h and low payoff ℓ, with 50% probability each. The FOSD lottery yields an additional premium $\pi > 0$. We implement three sets of parameters, $(\ell, h) \in \{(0, 180); (40, 140); (70, 110)\}$, and cross them with two different levels of the premium, $\pi \in \{4, 30\}$, resulting in six parameter combinations. Participants encounter each of the six tasks under three correlation structures: state-wise dominance, independence, and perfect negative correlation. The lottery pairs are designed judiciously to be simple and easily understood by participants. We use state-wise dominance and perfect negative correlation as these are the most extreme cases. Including the intermediate case of independence allows for a more severe test of the model.

Choice Display. Lottery pairs are described as "options" that yield payoffs depending on the turn of a wheel of fortune. The wheel implements the correlation structure in a natural and intuitive way. The lotteries are neutrally referred to as "Option A" and "Option B". We display lottery pairs in tabular form and as having four states to keep the displayed number of states constant across correlation structures, as illustrated in Figure 4, Panels (a)-(c). Participants are asked to choose the probability with which the computer implements either option by moving a slider to indicate the desired probabilities. To avoid anchoring effects, the initial position of the slider is hidden behind a blue bar as in (Grossmann, 2023). Once participants click on the bar, the slider appears. To make the choice interface more intuitive, the two sides of the slider are color coded to match the colors in which Options A and B are displayed, as highlighted in Figure 4, Panels (d) and (e).

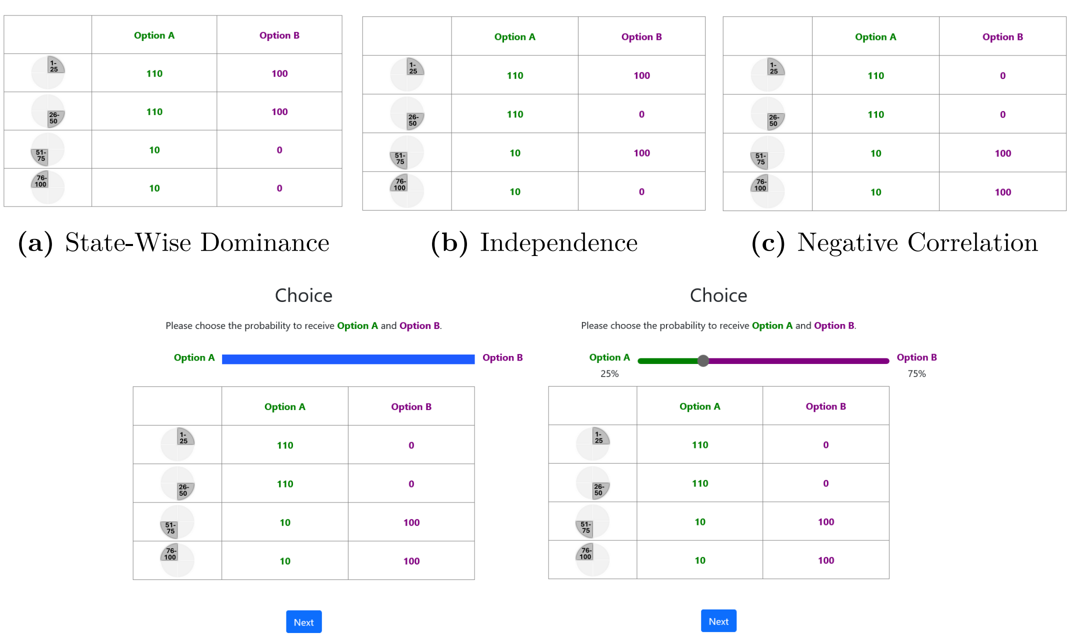

(a) State-Wise Dominance **(b)** Independence **(c)** Negative Correlation

(d) *Before* Clicking on the Slider **(e)** *After* Clicking on the Slider

Figure 4 Sample Screen Shots of the User Interface

Order and Practice Rounds. Participants encounter the different correlation structures in three distinct blocks of choices. Each block starts with four practice choices, for which participants receive immediate outcome feedback. Participants complete the practice choices for an additional choice task that is not payoff relevant. Each possible state occurs once, that is, the outcome distribution is representative. The practice rounds allow participants to familiarize themselves with the choice tasks. Within each block, the position of the dominant lottery and whether it is called Option A or Option B, as well as the order in which states are displayed, is fixed for each participant. Between blocks and between participants, the state order and which option is called A or B are randomized. Moreover, the order in which participants encounter the correlation structures is counterbalanced.

Prolific. We run the experiment on Prolific and implement a number of measures to ensure high data quality. In order to ensure language proficiency, we restrict access to the experiment to participants residing in the US or the UK. The instructions are followed by a set of three comprehension questions. If a participant answers a question incorrectly, they receive feedback designed to help them better understand the experiment. Participants who fail to answer the three comprehension questions correctly on the second attempt are screened out of the experiment. We also include three captchas at the beginning and require participants to complete the experiment on a computer rather than a phone or tablet to control how the tasks are presented.[6]

Attention and Sanity Tests. In addition to the choice tasks discussed above, each block contains a "sanity check". For this test, $(\ell, h, \pi) = (0, 1, 100)$. Participants always encounter this sanity test as the last choice in a block. They also face one attention test in each block, always between the second and third choice after the practice rounds. Attention tests feature tables similar to those used to present choice tasks but they displayed letters instead of numbers. Par-

[6] We ran the experiment before ChatGPT-5 was rolled out.

ticipants are asked to choose Option A with 56%.[7]

Payoffs. Participants received a fixed fee of £2.20. They took around 22 minutes to complete the experiment. In addition, one of their choices was randomly selected and paid out. Payoffs were displayed in pence and were equivalent to those given above. The average payment was £3.50, which corresponds to an average hourly wage of £9.90, and is 59% above the minimum hourly payment on Prolific.

We collected a total of 461 observations. In what follows, we restrict attention to our core sample of 228 participants who passed all three sanity tests, defined as selecting the FOSD lottery with $\gamma \geq 0.9$. Our results are robust to considering the full sample, but choice patterns become naturally more noisy.[8] The core sample is 43% female and 57% male, 42 years old on average, and mostly full- or part-time employed, with 10% being university students. The experiment was pre-registered and programmed using oTree (Chen et al., 2016).[9]

Discussion. A few aspects of the experimental design merit discussion. Different methods have been used to elicit PfR, such as repeated choice or introducing an option to flip a coin (Agranov and Ortoleva, 2017), as well as explicitly choosing the desired mixture (Feldman and Rehbeck, 2022; Agranov and Ortoleva, 2025). Our identification strategy relies on comparing rates of randomization across correlation structures. We therefore let participants choose their desired mixture, as this yields a fine-grained measure of PfR at the individual level. Asking participants to choose a mixture from a more restricted choice set, such as $\gamma \in \{0, 1\}$ or $\gamma \in \{0, 0.5, 1\}$, we still would be able to test our model at the aggregate level, but we might not be able to detect meaningful heterogeneity. Finally, we use sliders because other methods such as multiple choice lists have been shown to induce participants to avoid extremes (Beauchamp et al., 2020).

[7] Screenshots of the experimental interface can be found in the supplemental appendix.
[8] For details, see the supplemental appendix.
[9] See AEA Social Science Registry, ID AEARCTR-0015997.

We choose to vary the correlation of payoffs within participants, as this increase statistical power and allows us to study heterogeneity in PfR. However, our design also enables us to test our main hypotheses between-subjects by focusing on the first block that participants encounter. Participants are not informed in advance that they will encounter different correlation structures in the different blocks. They are only told that they have to make a number of choices that are grouped in three blocks. As we show in the supplemental appendix, we obtain very similar results if we compare behavior across participants.

Finally, each block of choices starts with four practice rounds. We made this design choice to improve participants' task understanding, ensure familiarity with the choice interface, and to let participants experience the consequences of their choices. Importantly, by exposing them to a representative sample of all possible states of the world, we ensure that participants experience the full joint distribution of payoffs.

4.2 Results: More Randomization under Less Concordant Correlation

Before turning to our main hypothesis, we first present the main summary statistics. Averaged over all three correlation structures, participants choose the dominant lottery with 80.2% probability. With 41.3% of all choices, $\gamma = 1$ is the modal choice, followed by $\gamma = 0.5$, which is chosen 9.0% of the time. 94.2% of all choices assign at least 50% weight to the FOSD lottery. Moreover, choice probabilities increase when the dominant lottery becomes relatively more attractive. Moving from $\pi = 4$ to $\pi = 30$, the probability with which participants choose the FOSD lottery increases from 76.1% to 84.3% ($p < 0.001$, OLS regression, standard errors clustered at the individual level).

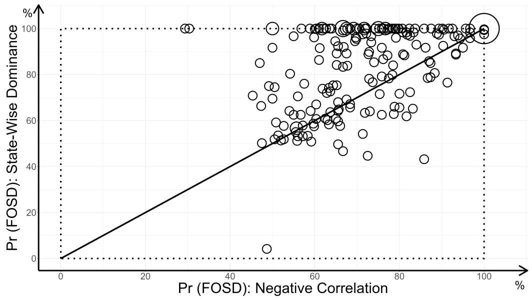

(a) State-Wise Dominance vs. Negative Correlation

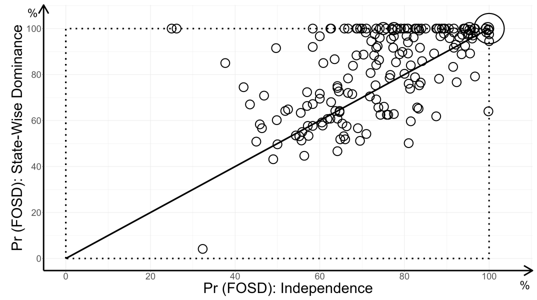

(b) State-Wise Dominance vs. Independence

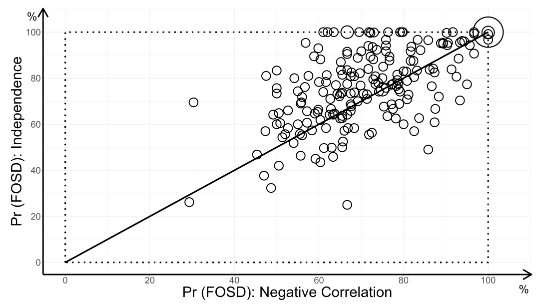

(c) Independence vs. Negative Correlation

Figure 5 Main Results: Randomization under Different Correlation Structures

In Figure 5, we present the results for our main hypothesis. As preregistered, we average choice probabilities by participant and correlation structure. On average, participants choose the dominant lottery with 85.9% probability under state-wise dominance, with 79.6% under independence, and with 75.1% probability under perfect negative correlation. All three pairwise comparisons are highly statistically significant with $p < 0.001$ (Wilcoxon signed-rank test). We thus confirm our main prediction: participants indeed randomize more under more negative correlation. In addition, participants are most likely to choose the FOSD lottery with probability one under state-wise dominance: 40.8% choose $\gamma = 1$ in this case, compared to only 20.6% under independence and 16.7% under negative correlation. Both differences are statistically significant at $p < 0.001$ (test of proportions). If we allow for a small margin of error of 5 percentage points, the fractions are 52.6%, 28.9%, and 18.9%, for state-wise dominance, independence, and negative correlation, respectively. The data thus also supports the prediction that participants are less likely to randomize at all when one lottery dominates states-wise.

In the supplemental appendix, we show that these findings are not driven by averaging choice probabilities by participant and correlation structure. We also obtain similar results when we consider the median instead of the mean, or when we simply consider all six choices a participant makes for a given correlation structure separately. We also use regression analyses to show that our results are robust to controlling for the order in which participants encounter different correlation structures, as well as round, block, and individual fixed effects.

Result 1. *Participants randomize more under more negative correlation structures. They choose the FOSD lottery with 85.9% under state-wise dominance, 79.6% under independence, and 76.1% under perfect negative correlation.*

4.3 Empirical Evidence of Ex-Post Regret

Our results so far confirm our main behavioral prediction that participants randomize more under more negative correlation. Next, we focus on the underlying mechanism. In our model, randomization is used to hedge anticipated regret-rejoice risk. The crucial feature necessary for correlation to impact randomization is that ex-post regret-rejoice is a function not just of the obtained, but also the forgone payoff. We next present survey evidence to test this feature of the model.

After making their decisions, participants are confronted with four different hypothetical scenarios in which they make choices similar to those in the main experiment, receive outcome feedback, and are asked to state their happiness. The scenarios comprise the independent correlation structure with parameters $(\ell, h, \pi) = (0, 100, 10)$. In all scenarios, participants are told that they chose the

FOSD lottery with 100% probability, in order to create clear predictions.[10] They are then asked to rate how they feel in these scenarios, on a nine-item Likert scale developed by Bradley and Lang (1994).[11] We code the scale such that the midpoint is normalized to 0, and options to the right (left) range from 1 (-1) to 4 (-4).

This setup allows an outside observer to test whether self-reported ex-post happiness respond to the obtained as well as the forgone outcome, as assumed in our model. We use independent correlation because it effectively creates a 2×2 design, where we cross participants' obtained payoff $\in \{10, 110\}$ with their forgone payoff $\in \{0, 100\}$. Thus, we can test whether self-reported happiness increases in the obtained payoff, while holding constant the forgone payoff, and vice versa. Testing these aspects of the model requires only that participants report happiness truthfully, and does not require cardinality. While self-reported happiness is not incentivized, participants have no incentive to misreport.

As predicted by the model, we find that self-reported happiness is increasing in the obtained, and decreasing in the forgone payoff. In the following, we denote $H(x, y)$ the reported happiness when obtaining outcome x and forgoing outcome y. We find that $H(110, 0) = 2.93 > H(110, 100) = 2.68 > H(10, 0) = 0.24 > H(10, 100) = -1.77$, with all possible comparisons being statistically significant at $p < 0.001$ (Wilcoxon signed-rank tests). Considering differences in happiness arising from different levels of the forgone outcome, the difference between $H(110, 0)$ and $H(110, 100)$ is rather small, whereas the difference between $H(10, 0)$ and $H(10, 100)$ is much larger. In the light of our model, this can be explained as follows. In the first comparison, participants remain in the rejoice domain, that is, in both cases the obtained payoff is higher than the forgone one. By contrast, in the second comparison, participants move from the rejoice

[10] When varying the mixture probability while holding outcomes constant, the predictions of our model depend on the participant's degree of OB λ. Therefore, we focus on the starkest case where $\gamma = 1$.

[11] For details, see the supplemental appendix.

to the regret domain when the forgone payoff increases from 0 to 100. Overall, self-reported happiness aligns with the predictions of our model. We summarize our findings below:

Result 2. *Self-reported ex-post happiness in hypothetical outcome scenarios increases in the obtained and decreases in the forgone payoff.*

4.4 A Typology of Randomizers

To explore heterogeneity in behavior, we next provide a typology of randomizers. Rather than analyzing behavior through the lens of our model, we let the data speak; that is, instead of imposing our model exogenously, we adopt a structure-free approach. We perform a clustering exercise using the k-means algorithm in the $[0, 100]^3$ space. The algorithm minimizes the within-cluster sum of squared distances and does not require imposing a specific functional form, which would be necessary if we were to perform structural estimation. In this sense, the k-means clustering algorithm is non-parametric. An analyst only has to specify the number of classes that they wish to obtain. In the supplemental appendix, we report elbow, proportional reduction in error, and silhouette analyses that suggest that the optimal number of classes is $k = 3$.[12] This number coincides with the three relevant model classes that are illustrated in Figure 3.

[12] Our analysis suggests that the increase in model fit decreases sharply after the third class is added. The classification we obtain is robust to different initial conditions. We perform 5,000 estimations with different initial condition. Only 36 of these yield a total of three unique classifications that differ from the one we present below. They all have a higher within-cluster sum of squares.

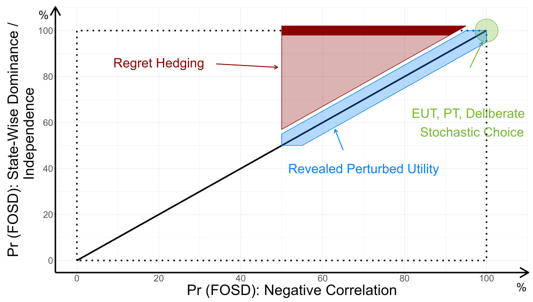

(a) Theoretical Predictions

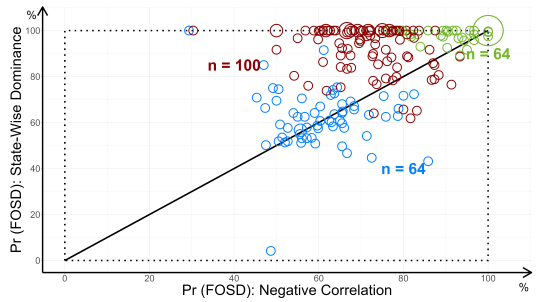

(b) State-Wise Dominance vs. Negative Correlation

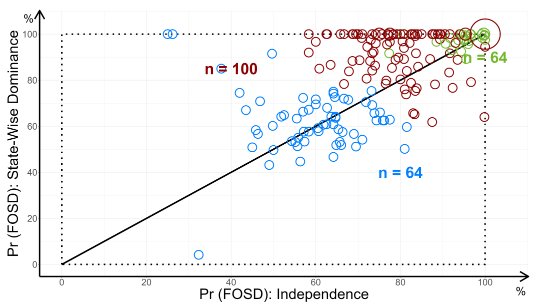

(c) State-Wise Dominance vs. Independence

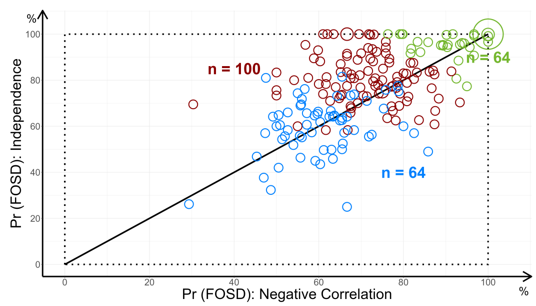

(d) Independence vs. Negative Correlation

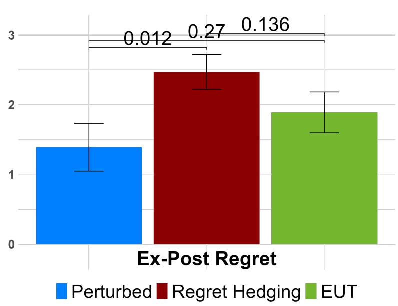

(e) Independence vs. Negative Correlation

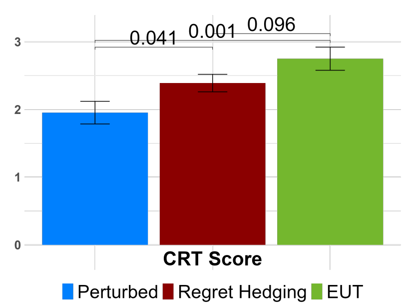

(f) Independence vs. Negative Correlation

Figure 6 Results of the k-Mean Clustering Exercise and Survey Responses

Panels (a)-(d) of Figure 6 present the results of this clustering exercise, along with the graphical illustration of the predictions of different models. Although the k-means algorithm does not impose any structural form, the three classes we obtain approximate the three theoretical classes surprisingly well. The largest class with 100 of the 228 participants displays PfR that depend strongly on the correlation structure. Participants in this class choose the FOSD lottery with 92.5% probability under state-wise dominance, with 80.7% under independence, and with 71.3% under negative correlation. We name this group the "regret-hedging class", as it is responsible for nearly all the aggregate-level correlation effects.[13]

The other two classes comprise 64 individuals each and display considerably less correlation sensitivity. Individuals in the second class choose the dominant lottery with high probabilities: 98.9% under state-wise dominance, 97.8% under independence, and 95.4% under negative correlation. We coin this the "EUT class". The final class exhibits a significantly lower propensity to choose the FOSD lottery: 62.4% under state-wise dominance, 59.7% under independence, and 60.8% under negative correlation. We refer to this as the "perturbed utility class".

Apart from providing a typology of randomizers, the clustering exercise also allows us to link survey responses to behavior, thereby providing further evidence on the mechanism underlying randomization.

4.5 The Correlates of Regret-Hedging

Anticipated Regret. If randomization is driven by anticipated regret-rejoice, participants who are assigned to the regret-hedging class should display the strongest response to forgone outcomes. We construct a measure of ex-post regret by taking the difference between $H(10, 0)$ and $H(10, 100)$. This measure captures the difference in reported happiness due to a change in the forgone outcome, holding the obtained outcome constant. As discussed above, the change in the counterfactual moves participants from the rejoice to the regret domain, which is a particularly salient comparison. Figure 6, Panels (e) and (f), illustrate

[13] The regret-hedging class exhibits significantly greater correlation-sensitivity in PfR than the other two classes. For each participant, we compute the difference between the probability of choosing the FOSD lottery under state-wise dominance and under negative correlation (the two alternative measures yield similar results) and test for differences across classes using Wilcoxon tests. The regret-hedging class is significantly more correlation-sensitive than either of the other classes ($p < 0.001$). Moreover, the EUT class shows greater correlation-sensitivity than the perturbed utility class ($p = 0.045$).

the results. Participants assigned to the regret-hedging class indeed have the highest average measure of ex-post regret. This measure is significantly higher compared to the perturbed utility class ($p = 0.012$, Wilcoxon rank-sum test) but fails to reach statistical significance compared to the EUT class ($p = 0.136$).

Cognitive Reflection. Incorporating anticipated ex-post regret into randomization decisions requires a certain level of sophistication. If participants' reactance to the correlation structure were driven by cognitive distortions, participants in the regret-hedging class should display low levels of sophistication. We measure sophistication using a four-item cognitive reflection test (CRT) proposed by Toplak et al. (2014). This test is inspired by the original CRT proposed by Frederick (2005) but uses different questions to address the concern that many participants may already have been exposed to the original CRT. The test is designed to measure the extent to which participants engage in deliberate, analytical thinking, and it is strongly correlated with measures of intelligence (Frederick, 2005; Toplak et al., 2014). We find that regret-hedgers have significantly higher CRT scores than their counterparts in the perturbed utility class ($p = 0.041$, Wilcoxon rank-sum test) and marginally significantly lower CRT scores ($p = 0.096$) than those participants who are assigned to the EUT class.

In sum, we find that regret-hedgers score high in ex-post regret and relatively high in cognitive reflection. This correlational evidence supports the mechanisms postulated in the model: participants randomize at different rates under different correlation structures because they deliberately hedge against ex-post post regret.

5 Experiment 2: Controlling for Display Effects

A potential concern about our findings presented so far is that correlation-sensitivity in PfR could be caused, not by deliberate regret-hedging, but aspects of the choice display. If participants compare payoffs row-by-row, it might be more challenging to see that one option is FOSD when payoffs are more negatively correlated. Two arguments speak against this concern. First, it is not obvious that independence is less cognitively demanding than the negative correlation. Independence features four distinct states, whereas the negative correlation has only two. Therefore, one might as well reach the opposite conclusion. Second, the correlational evidence presented above suggests that this explanation is unlikely. If correlation effects arose from confusion, those participants who score low in cognitive reflection would drive these effects. As discussed in Section 4.5, this is not the case. Moreover, the correlational evidence directly supports the regret-hedging channel.

To further address this possible concern and rule out possible confounds experimentally, we ran a second study with the goal of replicating the correlation effects in PfR while carefully controlling for display effects, and reducing cognitive demands to a minimum. Figure 7 illustrates the new choice display.

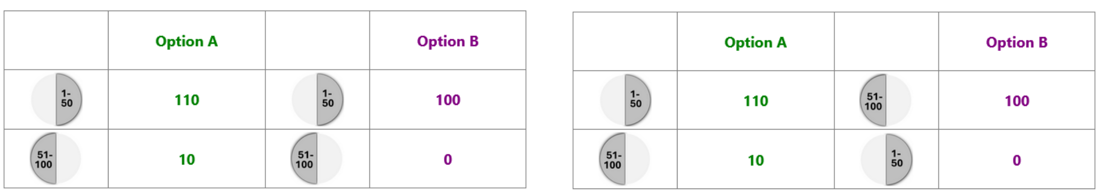

(a) State-Wise Dominance **(b)** Negative Correlation

Figure 7 Sample Screenshots from Experiment 2

Payoffs are now always perfectly aligned across rows, that is, the row-by-row comparison is held constant across correlation structures. We also reduce the number of displayed states from four to two in an additional effort to reduce complexity. As the independent correlation structure comprises four distinct states, it could be argued that it is inherently more demanding than state-wise dominance or negative correlation. Moreover, displaying independent lotteries in a way similar to the state-wise and negative correlation structures would introduce complications. For these reasons, we restrict the experiment to state-wise dominance and negative correlation. Apart from these changes, the design remains unchanged. That is, we use the same parameter values as before, and participants encounter each correlation structure in separate blocks of choices, each of which begins with four practice rounds.

We collected 189 observations on Prolific. Of these, 113 passed the sanity tests. In what follows, we focus on this sample. As before, results are robust to considering the full sample, as we show in the supplemental appendix. Averaging both correlation structures, participants choose the dominant lottery with 80.3% probability. The modal choice again is selecting the FOSD lottery with probability one (30.1% of choices), followed by $\gamma = 0.5$ (17.3% of choices). Furthermore, the probability of choosing the dominant lottery again increases with its premium. Moving from $\pi = 4$ to $\pi = 30$, the probability of choosing the dominant lottery increases from 76.4% to 84.2% ($p < 0.001$, OLS regression, standard errors clustered at the individual level). Overall, the choice probabilities of participants in this experiment are closely aligned to those in the main experiment.

The correlation effects persist even when we explicitly control for display effects. Figure 8 displays the main result.

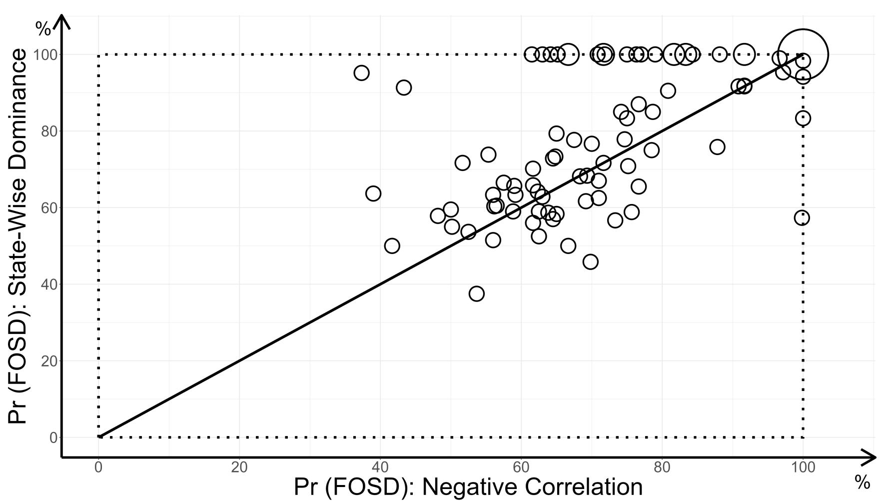

Figure 8 Randomization under Different Correlation Structures

On average, participants choose the FOSD lottery with 83.1% under state-wise dominance, and with 77.5% under negative correlation. Thus, while correlation effects are slightly reduced compared to the main experiment, the differences remain highly statistically significant ($p < 0.001$, Wilcoxon signed-rank test). Again, this result is robust to using the median mixture rather than the averaging by participant and correlation structure ($p < 0.001$, Wilcoxon signed-rank test). We view this as further evidence that the correlation effects in the main experiment are not caused by display effects, but are a manifestation of deliberate regret-hedging.

In addition, self-reported ex-post happiness again increases in the obtained and decreases in the forgone outcome, as predicted, although the difference between $H(100, 10)$ and $H(100, 100)$ now reaches statistical significance only at the 10% level ($p = 0.052$, Wilcoxon signed-rank). All other pairwise comparisons remain highly statistically significant at $p < 0.001$. We summarize our results below.

Result 3. *Participants randomize more under more negative correlation, even when carefully controlling for display effects.*

6 Experiment 3: Outcome Feedback

In this section, we examine the mechanism behind PfR more closely. If participants randomize to hedge against risk of anticipated regret from ex-post outcome comparison, DMs should adjust their behavior to the content of the outcome feedback they receive. In the literature on regret, the role of outcome feedback has been explored theoretically and its relevance for decision-making has been demonstrated experimentally (Bell, 1983; Zeelenberg et al., 1996; Engelbrecht-Wiggans and Katok, 2008; Filiz-Ozbay and Ozbay, 2007; Zheng, 2021; Dillenberger et al., 2025). If outcome feedback influences participants' randomization decisions, this would provide further support for the key role of regret-hedging.

We keep the setting as close as possible to the main experiment, and examine two distinct feedback structures: *Always Reveal* (AR) and *Choice Reveal* (CR). The former is the feedback structure we have considered thus far – the DM always receives feedback on the outcome of both lotteries, regardless of their choice. Under CR, the DM receives feedback only on the payoff they obtain, but not the non-chosen alternative. Both types of feedback structure occur naturally and are therefore important to understand. A prominent example of AR is the stock market, where investors observe the performance of all stocks, regardless of their individual investment choice. By contrast, CR captures situations where the non-chosen option never realizes. For example, if a manager decides against pursuing a risky project, this project is not undertaken, and hence the counterfactual outcome remains unknown. In what follows, we first derive theoretical predictions and then present experimental evidence.

6.1 Theoretical Predictions

We maintain the setup and the assumptions on the DM's utility function introduced above. However, the DM now forms expectations about the forgone outcome. As before, when observing the outcome, the DM engages in an ex-post evaluation of their choice, which might result in regret. For the implemented lottery A, we can still write the outcome-biased expected utility as $\lambda u(x_\omega^A) + (1-\lambda)EU^A$, as before. However, the DM now has to form an expectation about lottery B's outcome. We assume that they do so in an outcome-biased way. Upon observing that lottery A yields a payoff y, they inflate the probability of receiving this payoff. Conditional on this inflated probability, they form expectations in a Bayesian way.

Formally, denote the set of states in which lottery A yields a payoff y by $\Omega^y = \{\omega \in \Omega : x_\omega^A = y\}$, and the probability of lottery A yielding payoff y by $Pr(y) = \sum_{\omega\in\Omega^y} p_\omega$. The DM inflates this probability to $Pr(y) + \lambda\Big(1 - Pr(y)\Big)$ and assigns a weight $\frac{p_\omega}{Pr(y)}\Big[Pr(y) + \lambda\Big(1 - Pr(y)\Big)\Big]$ to state $\omega \in \Omega^y$. As before, they uniformly deflate the weights of the states $\omega \notin \Omega^y$ by (1-λ). We can write the outcome-biased expected utility of the non-implemented lottery B as:

$$\widetilde{EU}^{B}_{\omega\in\Omega^y} = \lambda E\Big[u(x^B_\omega)|y = x^A_\omega\Big] + (1-\lambda)EU^{\mathcal{L}}, \tag{1}$$

where $E\Big[u(x^B_\omega)|y = x^A_\omega\Big] = \sum_{j\in\Omega^y} u(x^{\mathcal{L}}_j)\frac{p_j}{Pr(y)}$. That is, for the non-implemented lottery B, the outcome-biased expected utility is a convex combination between the expectation over realized utilities, conditional on the observed payoff $y = x^A_\omega$, and the ex-ante expected utility, weighted by the outcome bias parameter λ. The regret-rejoicing experienced in state ω, when lottery $\mathcal{L}$ is implemented thus amounts to $R\Big((2\gamma-1)D^{\mathcal{L}}_\omega\Big)$, where $D^A_\omega = \lambda\Big(u(x^A_\omega) - E\Big[u(x^B_\omega)|y = x^A_\omega\Big]\Big) + (1-\lambda)\Delta_{EU}$. Likewise, $D^B_\omega = \lambda\Big(E\Big[u(x^A_\omega)|y = x^B_\omega\Big] - u(x^B_\omega)\Big) + (1-\lambda)\Delta_{EU}$.

The objective function becomes:

$$V_{CR}(\gamma) = \gamma\sum_{\omega\in\Omega} p_\omega R\Big((2\gamma-1)D^A_\omega\Big) + (1-\gamma)\sum_{\omega} p_\omega R\Big((2\gamma-1)D^B_\omega\Big). \tag{2}$$

Our formulation of the CR environment nests the AR case when the obtained outcome y is perfectly informative about the realized state. In this case, Ω^y is a singleton, $Pr(y) = p_\omega$ and the expression collapses to that used earlier. The outcome-biased decision weights sum up to one, and the formulation gives rise to a natural counterfactual – namely the expectation over the realized utilities of the non-chosen lottery. In previous work, this counterfactual is sometimes imposed exogenously (Gabillon, 2020; Zheng, 2021). In our model, this counterfactual follows naturally from the assumptions that regret ultimately results from information projection.

Comparing the DM's optimal mixture under the two feedback structures, we find that the CR environment introduces two distinct effects. While we build intuition for the main result here, we provide a more formal analysis in the supplemental appendix. First, partial outcome feedback makes ex-post comparisons less extreme, thereby reducing the regret-rejoice risk. This effect induces unambiguously less randomization under CR than AR. Second, there is an asymmetric feedback effect. In contrast to the AR environment, the feedback the DM receives now depends on their mixing probability γ. If the outcome feedback of one lottery shields the DM from ex-post regret, this lottery becomes relatively more attractive. For instance, if lottery A is a degenerate safe option, the DM effectively receives no outcome feedback when it is implemented, which eliminates regret-rejoice risk. Other than the regret-rejoice dampening effect of partial outcome feedback, the asymmetric feedback effect induces either more or less randomization. We derive clear predictions for our particular case.[14] As state-wise dominance and negative correlation are formally equivalent in both feedback structures, we focus on independence.

Prediction 2. *Consider two independent lotteries with marginal distributions* $(h+\pi, 0.5; \ell+\pi, 0.5)$ *and* $(h, 0.5; \ell, 0.5)$. *If* $u(\cdot)$ *is linear or concave, participants randomize less under the* CR *than under the* AR *feedback structure.*

We relegate the proof of Prediction 2 to the supplemental appendix.

6.2 Experimental Design

While accommodating the different feedback structures, we keep the design as consistent as possible with the main experiment, and focus on independence. The choice display is similar to that in the second experiment. That is, payoffs are displayed in two rows, and they are always perfectly aligned across rows. However, to implement independence, there are now two separate wheels of fortune: one determines the payoff of Option A, and the other one determines the payoff of Option B. We implement AR and CR as follows. Under AR, the computer always spins both wheels of fortune, and provides outcome feedback for both options.

[14] In the supplemental appendix, we discuss conditions under which CR induces unambiguously less randomization than AR.

Under CR, the computer spins only the wheel of fortune for the implemented option.

Except for these adjustments, the design remains unchanged. We implement the two feedback structures in a within-subject design, similar to how we implemented the different correlation structures in Experiments 1 and 2. That is, the choices are divided into two blocks of choices, one under AR and one under CR. As before, each block starts with four practice rounds that are not payoff relevant and for which participants receive immediate outcome feedback. Participants are not informed that they will encounter different feedback structures in different blocks, but receive instructions about the content of outcome feedback at the outset of each block. The parameters of the choice task are the same as in the first two experiments. Figure 9 illustrates the implementation.

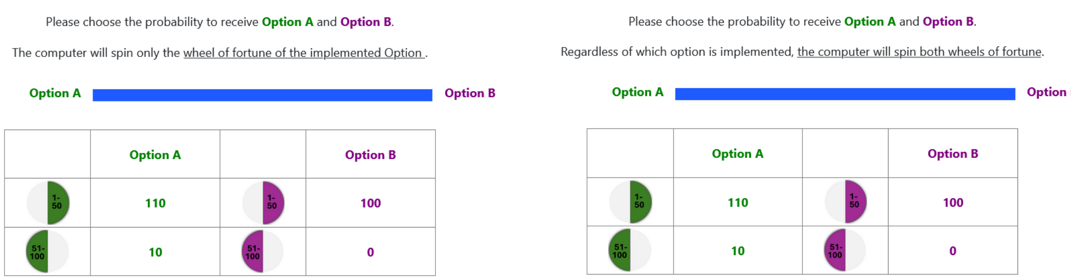

(a) Always Reveal (AR) Feedback Structure **(b)** Choice Reveal (CR) Feedback Structure

Figure 9 Sample Screenshots from Experiment 3

6.3 Findings

We recruited 196 participants on Prolific. In the following, we focus on the 117 participants who pass both sanity tests. As before, the results are robust to considering the whole sample. Averaging over both feedback structures, participants choose the FOSD lottery with 83.3% probability. The modal choice is again $\gamma = 1$ (51.1 % of choices), followed by $\gamma = 0.5$ (8.5% of choices). Participants again increase the probability with which they choose the dominant lottery when it becomes relatively more attractive. Moving from $\pi = 4$ to $\pi = 30$, the average probability of choosing the FOSD lottery increases from 79.8% to 86.8% ($p < 0.001$, OLS regression, standard errors clustered at the individual level).

We find that partial outcome feedback indeed reduces rates of randomization, as predicted by theory. Figure 10 displays the results.

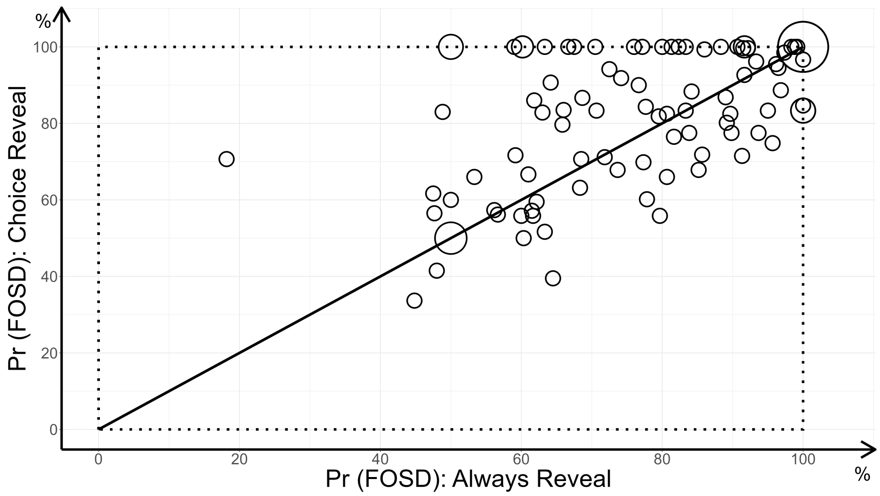

Figure 10 Randomization under Different Feedback Structures

Similar to previous figures, we plot the average mixture chosen by participants under AR against their average mixture under CR. The regret-hedging model predicts that the choices lie above the 45-degree line. Participants choose the dominant lottery with an average probability of 81.3% under AR, but with 85.3% under CR. The difference is statistically significant at $p = 0.035$ (Wilcoxon signed-rank test). Results remain similar when we use the median mixture rather than the average by participant and feedback structure ($p = 0.006$). Although feedback effects are arguably more subtle than the correlation effects we studied in experiments 1 and 2, participants adjusting their behavior to the feedback structure is consistent with recent evidence (Dillenberger et al., 2025).

Ex-Post Regret. As in Experiments 1 and 2, we elicit self-reported ex-post happiness after participants complete the main tasks. In addition to the four scenarios included in the previous experiments, we also add two scenarios in which participants receive only partial outcome information. In one scenario, they received the high payoff (110). In the other scenario, they received the low payoff (10). Considering the scenarios under AR, self-reported happiness again significantly increases with the obtained payoff, holding the forgone outcome constant. Holding the obtained outcome constant, we again find that self-reported happiness is significantly lower when the increase in the counterfactual moves participants from the rejoice to the regret domain ($p = 0.024$, Wilcoxon signed-rank test, $H(10, 0) = -1.08$, $H(10, 100) = -1.42$). When participants remain in the rejoice domain, we find no statistically significant difference when increasing the counterfactual from 0 to 100 ($H(110, 0) = 2.85$, $H(110, 100) = 3.11$, $p = 0.550$).

The regret hedging model predicts that partial outcome information renders regret-rejoice less extreme. The evidence from happiness reports is somewhat mixed. When the obtained outcome is low, we find $H(10, 0) = -1.08 < H(10, CR) = -1.39 < H(10, 100) = -1.42$. The first comparison is marginally statistically significant ($p = 0.099$), whereas the latter is not statistically significant ($p = 0.336$). When the obtained payoff is high, we find $H(110, CR) = 3.15 > H(110, 100) = 3.11 > H(110, 0) = 2.85$, which contradicts the predictions of the model. The difference between $H(110, CR)$ and $H(110, 100)$ is not statistically significant ($p = 0.284$), and the comparison between $H(110, CR)$ and $H(110, 10)$ is marginally statistically significant ($p = 0.070$).

7 Concluding Remarks

A growing body of literature documents preferences for randomization in decision-making under risk. When faced with a choice between two lotteries, many participants prefer to delegate the decision to a coin flip. In this paper, we propose a novel model in which individuals randomize to hedge against regret from unfavorable outcomes. With this model at hand, we can explain puzzling findings in the literature, such as high rates of randomization even when one option first-order stochastically dominates the other.

We conduct three experiments to test the model. In the first experiment, we test the key novel prediction that rates of randomization depend on how payoffs are correlated across states, and find that participants randomize more when lottery payoffs are more negatively correlated. We provide a typology of randomizers that separates regret-hedging from other models accommodating preferences for randomization, namely perturbed utility and deliberate random choice models. Regret-hedging emerges as the most prominent type. Participants who are classified as this type score high in self-reported ex-post regret and cognitive reflection. In the second experiment, we rule out potential confounds due to display effects. In the third experiment, we test an additional unique prediction of the model: the extent to which participants randomize depends on the content of outcome feedback. Although this prediction is arguably more subtle, we document evidence that withholding outcome feedback on the non-chosen lottery decreases rates of randomization. Overall, the evidence suggests that preferences for randomization are a manifestation of regret-hedging.

References

Agranov, M., P. J. Healy, and K. Nielsen (2023). Stable randomisation. *The Economic Journal 133*(655), 2553–2579.

Agranov, M. and P. Ortoleva (2017). Stochastic choice and preferences for randomization. *Journal of Political Economy 125*(1), 40–68.

Agranov, M. and P. Ortoleva (2022). Revealed preferences for randomization: An overview. In *AEA Papers and Proceedings*, Volume 112, pp. 426–430.

Agranov, M. and P. Ortoleva (2025). Ranges of randomization. *Review of Economics and Statistics*, 1–12.

Baron, J. and J. C. Hershey (1988). Outcome bias in decision evaluation. *Journal of Personality and Social Psychology 54*(4), 569.

Beauchamp, J. P., D. J. Benjamin, D. I. Laibson, and C. F. Chabris (2020). Measuring and controlling for the compromise effect when estimating risk preference parameters. *Experimental Economics 23*(4), 1069–1099.

Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research 30*(5), 961–981.

Bell, D. E. (1983). Risk premiums for decision regret. *Management Science 29*(10), 1156–1166.

Bordalo, P., N. Gennaioli, and A. Shleifer (2012). Salience theory of choice under risk. *The Quarterly Journal of Economics 127*(3), 1243–1285.

Bradley, M. M. and P. J. Lang (1994). Measuring emotion: The self-assessment manikin and the semantic differential. *Journal of Behavior Therapy and Experimental Psychiatry 25*(1), 49–59.

Camerer, C. F. (1989). An experimental test of several generalized utility theories. *Journal of Risk and uncertainty 2*(1), 61–104.

Cerreia-Vioglio, S., D. Dillenberger, and P. Ortoleva (2015). Cautious expected utility and the certainty effect. *Econometrica 83*(2), 693–728.

Cerreia-Vioglio, S., D. Dillenberger, P. Ortoleva, and G. Riella (2019). Deliberately stochastic. *American Economic Review 109*(7), 2425–2445.

Cettolin, E. and A. Riedl (2019). Revealed preferences under uncertainty: Incomplete preferences and preferences for randomization. *Journal of Economic Theory 181*, 547–585.

Chen, D. L., M. Schonger, and C. Wickens (2016). otree - an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance 9*, 88–97.

Chen, Y.-C., S. H. Chew, and X. Zhang (2025). Correlation preference. *Available at SSRN 5253868*.

Chew, S. H., L. G. Epstein, and U. Segal (1991). Mixture symmetry and quadratic utility. *Econometrica*, 139–163.

Chew, S. H., B. Miao, Q. Shen, and S. Zhong (2022). Multiple-switching behavior in choice-list elicitation of risk preference. *Journal of Economic Theory 204*, 105510.

Danz, D., D. Kübler, L. Mechtenberg, and J. Schmid (2015). On the failure of hindsight-biased principals to delegate optimally. *Management Science 61*(8), 1938–1958.

Diecidue, E. and J. Somasundaram (2017). Regret theory: A new foundation. *Journal of Economic Theory 172*, 88–119.

Dillenberger, D., Y. Halevy, J. Hoelzemann, G. Nave, and P. Rahmani (2025). Anticipated regret. Technical report, University of Pennsylvania.

Dwenger, N., D. Kübler, and G. Weizsäcker (2018). Flipping a coin: Evidence from university applications. *Journal of Public Economics 167*, 240–250.

Engelbrecht-Wiggans, R. and E. Katok (2008). Regret and feedback information in first-price sealed-bid auctions. *Management Science 54*(4), 808–819.

Feldman, P. and J. Rehbeck (2022). Revealing a preference for mixtures: An experimental study of risk. *Quantitative Economics 13*(2), 761–786.

Filiz-Ozbay, E. and E. Y. Ozbay (2007). Auctions with anticipated regret: Theory and experiment. *American Economic Review 97*(4), 1407–1418.

Fischhoff, B. (1975). Hindsight is not equal to foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance 1*(3), 288.

Fishburn, P. C. (1982). Nontransitive measurable utility. *Journal of Mathematical Psychology 26*(1), 31–67.

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives 19*(4), 25–42.

Fudenberg, D., R. Iijima, and T. Strzalecki (2015). Stochastic choice and revealed perturbed utility. *Econometrica 83*(6), 2371–2409.

Gabillon, E. (2020). When choosing is painful: Anticipated regret and psychological opportunity cost. *Journal of Economic Behavior & Organization 178*, 644–659.

Gollier, C. (2020). Aversion to risk of regret and preference for positively skewed risks. *Economic Theory 70*(4), 913–941.

Grossmann, M. (2023). Otree slider. Technical report.

Gul, F. (1991). A theory of disappointment aversion. *Econometrica*, 667–686.

Gurdal, M. Y., J. B. Miller, and A. Rustichini (2013). Why blame? *Journal of Political Economy 121*(6), 1205–1247.

Hey, J. D. (2001). Does repetition improve consistency? *Experimental Economics 4*, 5–54.

Hey, J. D. and C. Orme (1994). Investigating generalizations of expected utility theory using experimental data. *Econometrica*, 1291–1326.

Heydari, P. (2024). Regret, responsibility, and randomization: A theory of stochastic choice. *Journal of Economic Theory 217*, 105824.

Kahneman, D. and A. Tverski (1979). Prospect theory: An analysis of decision under risk. *Econometrica 47*(2), 263–292.

Lanzani, G. (2022). Correlation made simple: Applications to salience and regret theory. *Quarterly Journal of Economics 137*(2), 959–987.

Levitt, S. D. (2021). Heads or tails: The impact of a coin toss on major life decisions and subsequent happiness. *The Review of Economic Studies 88*(1), 378–405.

Loewenfeld, M. (2025). Outcome bias and delegated decision-making: Theory and experiment. Technical report, University of Vienna.

Loewenfeld, M. and J. Zheng (2024). Salience or event-splitting? an experimental investigation of correlation sensitivity in risk-taking. *Journal of the Economic Science Association 10*(2), 346–366.

Loewenfeld, M. and J. Zheng (2025). Uncovering correlation sensitivity in decision making under risk. *International Economic Review*. Forthcoming.

Loomes, G. (1998). Probabilities vs money: A test of some fundamental assumptions about rational decision making. *The Economic Journal 108*(447), 477–489.

Loomes, G. and R. Sugden (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal 92*(368), 805–824.

Loomes, G. and R. Sugden (1987). Some implications of a more general form of regret theory. *Journal of Economic Theory 41*(2), 270–287.

Loomes, G. and R. Sugden (1995). Incorporating a stochastic element into decision theories. *European Economic Review 39*(3-4), 641–648.

Machina, M. J. (1985). Stochastic choice functions generated from deterministic preferences over lotteries. *The Economic Journal 95*(379), 575–594.

Madarász, K. (2012). Information projection: Model and applications. *The Review of Economic Studies 79*(3), 961–985.

Manzini, P. and M. Mariotti (2014). Stochastic choice and consideration sets. *Econometrica 82*(3), 1153–1176.

Milgrom, P. and C. Shannon (1994). Monotone comparative statics. *Econometrica*, 157–180.

Milligan, G. W. and M. C. Cooper (1985). An examination of procedures for determining the number of clusters in a data set. *Psychometrika 50*(2), 159–179.

Mosteller, F. and P. Nogee (1951). An experimental measurement of utility. *Journal of political economy 59*(5), 371–404.

Müller, A. and D. Stoyan (2002). *Comparison Methods for Stochastic Models and Risks*. Chichester: Wiley.

Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior & Organization 3*(4), 323–343.

Quiggin, J. (1990). Stochastic dominance in regret theory. *The Review of Economic Studies 57*(3), 503–511.

Quiggin, J. (1994). Regret theory with general choice sets. *Journal of Risk and Uncertainty 8*, 153–165.

Rousseeuw, P. J. (1987). Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics 20*, 53–65.

Rubinstein, A. (2002). Irrational diversification in multiple decision problems. *European Economic Review 46*(8), 1369–1378.

Sarver, T. (2008). Anticipating regret: Why fewer options may be better. *Econometrica 76*(2), 263–305.

Savage, L. J. (1954). *The Foundations of Statistics*. New York: Courier Corporation.

Sopher and Narramore (2000). Stochastic choice and consistency in decision making under risk: An experimental study. *Theory and Decision 48*(4), 323–349.

Thorndike, R. L. (1953). Who belongs in the family? *Psychometrika 18*(4), 267–276.

Topkis, D. M. (1978). Minimizing a submodular function on a lattice. *Operations Research 26*(2), 305–321.

Toplak, M. E., R. F. West, and K. E. Stanovich (2014). Assessing miserly information processing: An expansion of the cognitive reflection test. *Thinking & Reasoning 20*(2), 147–168.

Toussaert, S. (2024). Stochastic dominance and demand for surprise. Technical report, University of Oxford.

Tversky, A. (1969). Intransitivity of preferences. *Psychological Review 76*(1), 31–48.

Tversky, A. and D. Kahneman (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty 5*(4), 297–323.

Zeelenberg, M., J. Beattie, J. Van der Pligt, and N. K. De Vries (1996). Consequences of regret aversion: Effects of expected feedback on risky decision making. *Organizational Behavior and Human Decision Processes 65*(2), 148–158.

Zeelenberg, M., W. W. Van Dijk, A. S. Manstead, and J. vanr de Pligt (2000). On bad decisions and disconfirmed expectancies: The psychology of regret and disappointment. *Cognition & Emotion 14*(4), 521–541.

Zheng, J. (2021). Willingness to pay for reductions in health risks under anticipated regret. *Journal of Health Economics 78*, 102476.