Do Hedgers Minimize Variance or Downside Risk?

Paul Feldman
Texas A&M University

Siun Lee
Texas A&M University

Dmitry Vedenov
Texas A&M University

Abstract

Traditional hedging models recommend minimizing the variance of returns, treating upside and downside risk symmetrically. This approach contradicts both economic intuition and behavioral evidence showing that decision-makers care more about potential losses than equivalent gains. We conduct controlled experiments to test whether participants' hedging decisions align more closely with Minimum Variance (MV) or Lower Partial Moment LPM_2 optimization, where LPM_2 focuses exclusively on downside risk. Our results reveal substantial heterogeneity in hedging behavior, with a significant proportion of participants better described by LPM_2 than MV criteria. The superiority of LPM_2 is most pronounced in scenarios with asymmetric risk profiles, where MV's symmetric treatment systematically underestimates participants' loss aversion. These patterns align with predictions from cumulative prospect theory, particularly the tendency to overweight losses relative to gains and to emphasize small-probability extreme outcomes. Our findings challenge the widespread use of variance-based measures in risk management and suggest that models incorporating asymmetric risk preferences provide more accurate predictions of hedging behavior. The results have important implications for the design of hedging strategies and risk management frameworks that better reflect how decision-makers actually perceive and respond to risk.

Keywords: hedging, downside risk, lower partial moment, minimum variance, prospect theory, experimental economics *JEL Classification:* G11, G13, D81, C91

1 Introduction

Consider a farmer hedging next year's corn harvest using futures contracts. The key question she needs to answer is how many futures contracts to use, or alternatively, what *hedge ratio* to choose[1]. Traditional hedging models typically recommend selecting hedge ratios by minimizing the total variance of profits thus treating a $10,000 unexpected gain equally as "risky" as a $10,000 unexpected loss. This approach, known as minimum variance (MV) strategy, is widely used in the hedging literature. However, it contradicts basic economic intuition: the farmer clearly cares more about potential losses that threaten financial survival rather than gains that exceed expectations.

[1]Hedge ratio is defined as a ratio of the size of the futures position to the size of the asset being hedged.

Variance-based measures become particularly problematic when return distributions exhibit skewness or heavy tails — features common in commodity markets. MV hedge ratios may systematically over-hedge thus reducing beneficial upside volatility or under-hedge thus failing to adequately protect against extreme losses. This symmetric treatment of risk contradicts extensive behavioral evidence showing that individuals are loss averse and place greater weight on avoiding losses rather than on realizing equivalent gains (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992).

An alternative approach — minimization of the downside risk, commonly measured by the second-order lower partial moment (LPM_2), — aligns more closely with how hedgers perceive risk. Unlike variance, which penalizes both gains and losses, LPM_2 captures only downside risk—specifically, the expected squared deviations below a target return. In our farmer's case, LPM_2 focuses exclusively on the outcomes that fall short of expectations, ignoring upside fluctuations that are generally welcome.

Despite the theoretical appeal of the downside risk measures, empirical evidence on whether actual hedging decisions reflect MV or LPM_2 optimization remains limited. While several studies have applied LPM_2 in agricultural and commodity contexts (e.g., Mattos, Garcia, and Nelson, 2008; Power and Vedenov, 2010; Liu et al., 2017), a fundamental question remains: Do real-world decision-makers hedge to minimize variance, or do they hedge to minimize downside risk?

This paper addresses this question through a series of controlled experiments on hedging behavior. We test whether participants' hedging choices align more closely with the MV or LPM_2-based optimal hedge ratios.

Understanding which risk measure better describes hedging behavior has important implications for theory and practice. If downside risk measures provide a better fit, this would support incorporating behavioral insights—such as loss aversion—into risk management models. More accurate models could, in turn, inform the design of hedging strategies that better reflect how decision-makers perceive and respond to risk, potentially improving hedging effectiveness and adoption.

While LPM_2 can be justified on normative grounds, it also reflects behavioral tendencies predicted by prospect theory. Loss aversion leads decision makers to focus on avoiding returns below a reference point, while probability weighting causes them to overemphasize small-probability catastrophic outcomes. Mattos, Garcia, and Nelson (2008) show that probability weighting is the dominant behavioral factor influencing optimal hedge ratios, with loss aversion and risk aversion playing smaller roles unless combined with probability weighting.

Despite these insights, little is known about whether actual hedging decisions align more closely with MV, LPM_2, or are systematically shaped by prospect-theoretic parameters. This paper addresses this gap through a series of controlled experiments on hedging behavior, testing whether participants' choices better match the MV or LPM_2 optimal hedge ratios and whether deviations from either can be explained by loss aversion and probability weighting.

Our study contributes to two key areas of research. First, it provides empirical evidence comparing the descriptive validity of different risk measures in explaining hedging behavior. Second, it examines whether prospect theory can account for systematic deviations from traditional models, highlighting the role of behavioral biases in risk management decisions.

In section 3, literature review. Section 4, It briefly explains hedging problem of the real world. Section 5. Experimental design. And REsult. And then we conclude.

2 Literature Review

The choice between minimum variance (MV) and downside risk minimization remains a fundamental debate in hedging theory. While MV optimization has dominated practice since the 1960s, growing evidence supports decision-makers' preference for second-order lower partial moment (LPM_2) approaches that focus on downside protection. This review examines the theoretical foundations and empirical evidence for both frameworks, identifying critical gaps in experimental validation.

Johnson (1960)[Joh60] established variance minimization as the foundation for optimal hedging, with Ederington (1979)[Ede79] rationalizing the framework through the now-standard hedge ratio formula $h^* = \frac{Cov(R_s,R_f)}{Var(R_f)}$. Rooted in Markowitz's (1952)[Mar52] mean-variance portfolio theory, this approach offers mathematical tractability but rests on restrictive assumptions—quadratic utility, symmetric risk preferences, and normally distributed returns. Most critically, variance penalizes upside and downside deviations equally, contradicting extensive evidence that decision-makers care more about avoiding losses than forgoing equivalent gains.

Recognition of these limitations led to downside risk measures rooted in Roy's (1952) safety-first principle. Bawa (1975) formalized Lower Partial Moments as $LPM_n(\tau) = E[max(\tau - R, 0)^n]$, with LPM_2 capturing semivariance through quadratic penalties for returns below target return τ. Markowitz himself later acknowledged that semivariance is the "more plausible" risk measure. LPM_2 accommodates non-normal distributions, focuses on losses, and is consistent with second-degree stochastic dominance.

Empirical studies demonstrate LPM_2's advantages over MV in hedging applications. Mattos, Garcia, and Nelson (2008)[MGN08] find that LPM_2 hedge ratios can diverge substantially from MV ratios under asymmetric distributions. Power and Vedenov (2010)[PV10] show that LPM_2 strategies typically yield smaller hedge ratios but achieve superior downside risk control, challenging the assumption that greater hedging intensity automatically improves risk management. Liu, Vedenov, and Power (2017) [LVP17] extend these results, showing that LPM_2 strategies deliver higher expected profits and lower expected shortfall, particularly under changing dependence structures among commodity prices.

While MV and LPM_2 describe the objective distribution of outcomes, they do not account for how individuals perceive and value those outcomes. Prospect Theory (Kahneman & Tversky, 1979[TK79]; Tversky & Kahneman, 1992[TK92]) provides a behavioral framework that can bridge this gap. Its two core elements—loss aversion and probability weighting—predict that decision-makers overweight small-probability disasters and underweight moderate-probability events, amplifying the salience of tail risk.

Mattos, Garcia, and Pennings (2008)[MGP08] show that probability weighting can shift optimal hedge ratios toward values resembling LPM_2 prescriptions, while loss aversion alone has little impact without such weighting. Survey evidence reinforces this behavioral tendency: Deloitte reports that 78% of CFOs explicitly hedge to minimize downside risk, not variance. Together, these findings suggest that observed hedging choices may deviate from both MV and LPM_2 optima due to systematic behavioral distortions.

Despite extensive theoretical and empirical support for LPM_2, there is a critical lack of controlled experimental evidence testing whether hedgers' decisions align more with MV, LPM_2 or Prospect Theory–adjusted prescriptions. Most existing work relies on historical data or surveys, leaving open whether observed patterns reflect genuine preferences, cognitive biases, or institutional constraints. Experimental economics research

highlights related challenges: Blanco et al. (2010) note that hedging opportunities can confound belief elicitation, and identity signaling studies find that individuals sometimes reject "free" hedging opportunities. Furthermore, almost no studies use real-money experiments, limiting ecological validity, and cross-cultural variation in risk measure preferences remains unexplored.

Implementation barriers also persist. LPM_2 estimation requires more historical data and lacks standardized backtesting frameworks, while regulatory and accounting standards remain variance-based (e.g., 80–125% effectiveness thresholds). Evidence from crisis periods such as the COVID-19 pandemic shows that optimal hedge ratios may vary dramatically across regimes, suggesting that model choice may depend on market stress conditions.

3 Hedging Problem

In a typical hedging decision, a seller faces uncertainty about future prices. At the start (t=0), they know the current spot price (for immediate delivery) and the current futures price (for delivery in the future). The actual transaction will occur later (t=1), when both prices may have changed.

If the seller does not hedge, their revenue depends entirely on the spot price at the time of sale:

$$\pi_{\text{unhedged}} = S_1 = S_0 + \Delta S,$$

where ΔS = S_1 - S_0 is the change in the spot price. A positive ΔS means a gain, while a negative ΔS means a loss.

Alternatively, the seller can take a short position in the futures market to lock in part of the price in advance. This creates two payoff components: 1) the spot market revenue, which will still fluctuate with the spot price, and 2) the futures market payoff, which typically move in the opposite direction to the spot price.

In the experiment, we model the spot and futures positions as perfectly negatively correlated lotteries for simplicity, even though in real markets spot and futures prices are typically highly positively correlated.

Let the hedge ratio be

$$h := \frac{Q_F}{Q_A},$$

measures the size of the futures position relative to the spot position. In per-unit terms, the hedged payoff is:

$$\pi(h) \;=\; S_1 + h(F_0 - F_1) \;=\; S_0 + \big(\Delta S - h\,\Delta F\big), \tag{1}$$

where ΔF = F_1 - F_0 is the change in the futures price. By adjusting the hedge ratio—the proportion of their position covered in the futures market—the seller can manage the trade-off between risk and return.

In theory, the "optimal" hedge ratio depends on how we define and measure risk. Two common approaches are:

Minimum Variance (MV) – focuses on reducing overall price variability. Under this approach, the optimal hedge ratio balances movements in the spot and futures prices

so that gains in one market offset losses in the other as much as possible. A variance minimizer solves

$$h^{MV} \in \arg\min_h \mathrm{Var}\big(\pi(h)\big) \equiv \arg\min_h \mathrm{Var}\big(\Delta S - h\,\Delta F\big),$$

which yields the standard solution

$$h^{MV} = \frac{\mathrm{Cov}(\Delta S, \Delta F)}{\mathrm{Var}(\Delta F)} = \rho_{SF}\,\frac{\sigma_S}{\sigma_F}.$$

When $\rho_{SF} \approx 1$, this reduces approximately to $h^{MV} \approx \sigma_S/\sigma_F$.

LPM_2 focuses on downside risk — the risk of earning less than a target revenue. Rather than reducing all variability, it penalizes large shortfalls more heavily. This approach reflects a more loss-averse perspective.

$$h^{LPM_2} \in \arg\min_h LPM_2(h;\bar{\pi}) = \mathbb{E}\Big[\big(\bar{\pi} - \pi(h)\big)_+^2\Big] = \int_{-\infty}^{\bar{\pi}} \big(\bar{\pi} - \pi\big)^2\, dF_\pi(\pi),$$

where $(x)_+ := \max\{x, 0\}$ and F_π is the distribution of $\pi(h)$.

In our experiment, participants face a simplified version of this decision. Instead of real prices, they allocate between two perfectly negatively correlated lotteries representing the spot and futures positions. The structure mirrors the hedging problem: the allocation choice corresponds to setting a hedge ratio, and we can observe whether participants' behavior aligns more with MV reasoning or with LPM_2 reasoning.

4 Experimental Design

This experiment tests whether hedging decisions align more closely with symmetric risk measures (Minimum Variance, MV) or downside-focused measures (LPM_2). We use a controlled portfolio allocation task in which participants allocate between two perfectly negatively correlated lotteries, mimicking the relationship between spot and futures prices. This structure allows us to manipulate payoff asymmetry and variance independently, creating scenarios where MV and LPM_2 predict different optimal hedge ratios.

4.1 Task

In each round, participants allocate their money between two binary-outcome lotteries: L_1 (spot position) and L_2 (futures position). The lotteries are perfectly negatively correlated ($\rho = -1$), so a favorable outcome for one implies an unfavorable outcome for the other. In addition, the probability p shapes both the portfolio's expected payoff and risk, influencing how participants choose their allocations.

The outcomes are defined as follows:

Let $\alpha \in [0,1]$ denote the fraction invested in L_1 (with $1-\alpha$ in L_2). The portfolio payoff is:

$$\text{Outcome} = \begin{cases} \alpha x_{\text{low}} + (1-\alpha) y_{\text{high}}, & \text{State 1,} \\ \alpha x_{\text{high}} + (1-\alpha) y_{\text{low}}, & \text{State 2.} \end{cases}$$

Table 1: Payoff Matrix for Binary-Outcome Lotteries

State	**Probability**	L_1 **(Spot)**	L_2 **(Futures)**
1	p	x_{low}	y_{high}
2	$1-p$	x_{high}	y_{low}

The hedge ratio is:

$$h = \frac{1-\alpha}{\alpha},$$

where $h = 0$ indicates no hedging (full allocation to L_1), and larger values represent more aggressive hedging. This formulation allows for intuitive interpretation for α. If $\alpha = 1$ corresponds to no hedging (all funds in L1, $h = 0$), while decreasing values of α indicate increasingly aggressive hedging strategies with higher hedge ratios.

4.2 Scenario Overview

Across five scenarios, we vary payoff structures ($x_{\text{low}}, x_{\text{high}}, y_{\text{low}}, y_{\text{high}}$) and probabilities ($p$) to manipulate symmetry, variance, and downside exposure. This design generates distinct optimal hedge ratios under the MV and LPM_2 frameworks, allowing us to compare theoretical predictions with actual participant decisions. Table 2 summarizes the key features and model predictions based on the actual experimental parameters.

Table 2: Scenario Overview

Scen.	**Symmetry**	**Equal Variance**	**Skewness**	**Probabilities**	**MV** α	LPM_2 α
1	Symmetric	Yes	None	0.5	0.5	0.5
2	Asymmetric	Yes	Moderate	Varies	0.5	Decreasing over rounds
3	Asymmetric	No (L2 higher)	High	Varies	≈ 0.76	Decreasing over rounds
4	Asymmetric	No (L2 higher)	Extreme	Varies	0.5	Decreasing over rounds
5	Asymmetric	Yes	Moderate	0.5	0.5	Decreasing over rounds

5 Result

5.1 Overall Model Performance

We analyze 45 decision rounds from 148 participants. Participant choices are predominantly above 0.5; across rounds, the mean of participant means is $\bar{m} = 0.66$ (median 0.66), and robust summaries (trimmed) are slightly higher. Model predictions are fixed by scenario and round: MV concentrates near $\{0.5, 0.76\}$ depending on scenario, while LPM_2 varies more widely within scenarios, especially in Scenarios 2 to 4.

Using the row mean of participant choices as the target, LPM_2 exhibits lower squared and absolute error than MV overall. The mean squared error (MSE) is 0.02 for LPM_2 versus 0.03 for MV; mean absolute error (MAE) is 0.14 for LPM_2 versus 0.15 for MV. A participant–bootstrap (resampling columns of the 45×148 matrix) confirms that the average performance difference favors LPM_2 but is small in magnitude. The bootstrap

mean MSEs are 0.03 for MV and 0.03 for LPM_2, with the mean difference (LPM_2 − MV) equal to −0.000964 and a 95% percentile interval [−0.002662, 0.000813], implying no statistically significant difference at the 5% level. Nevertheless, LPM_2 outperforms MV in 86.6% of bootstrap draws, which is consistent with a modest but directionally stable advantage.

The estimator–sensitivity checks (10%/20% trimmed means and 5–95% winsorized means) preserve the aggregate ranking. Scenario analyses show that Scenarios 2, 4, and 5 favor LPM_2, Scenario 1 is a tie, and Scenario 3 favors MV when models are judged against the round mean. To align inference with individual decisions, we complement these aggregates with a participant–level, paired–loss evaluation

$$\Delta = \left| y - \widehat{y}^{\mathrm{LPM}_2} \right| - \left| y - \widehat{y}^{\mathrm{MV}} \right|,$$

so that $\Delta < 0$ favors LPM_2 and $\Delta > 0$ favors MV. In Scenario 3 the distribution of Δ concentrates slightly to the right of zero, indicating that MV is closer to participants more often, albeit by small margins; in the other scenarios the paired–loss results are consistent with the aggregate ranking and reinforce LPM_2's advantage. Taken together, the evidence supports LPM_2 as the primary specification for this setting, with MV remaining a competitive benchmark in environments resembling Scenario 3.

Table 3: Overall performance using the row mean as the target

Model	MSE	MAE	Bootstrap MSE (mean)	95% CI (MSE)
MV	0.03	0.15	0.03	[0.02, 0.03]
LPM_2	0.02	0.14	0.03	[0.02, 0.03]

Beyond the aggregate, the estimator sensitivity check shows the same story. When the participant aggregate is computed as the mean, median, 10%/20% trimmed mean, or 5–95% winsorized mean, LPM_2 consistently attains the lower MSE than MV, with gaps ranging from about 0.00 to 0.01. Two quartile–based summaries illustrate the source of heterogeneity: using the lower quartile (q25) as the target, MV is better (MSE 0.026 vs. 0.043), indicating LPM_2 sits above the most conservative responders; using the upper quartile (q75), both models are less accurate, but LPM_2 remains better (MSE 0.17 vs. 0.20). These patterns are consistent with participant distributions centered well above 0.5 and with LPM_2 tracking the central tendency more closely than MV in most scenarios.

5.2 Scenario-Specific Analysis

Scenario-level results help clarify the sources of the aggregate performance differences between MV and LPM_2. Table 4 reports mean squared error (MSE) and mean absolute error (MAE) for each scenario, using the scenario–round mean of participant allocations as the target. Negative Δ values favor LPM_2, while positive values favor MV.

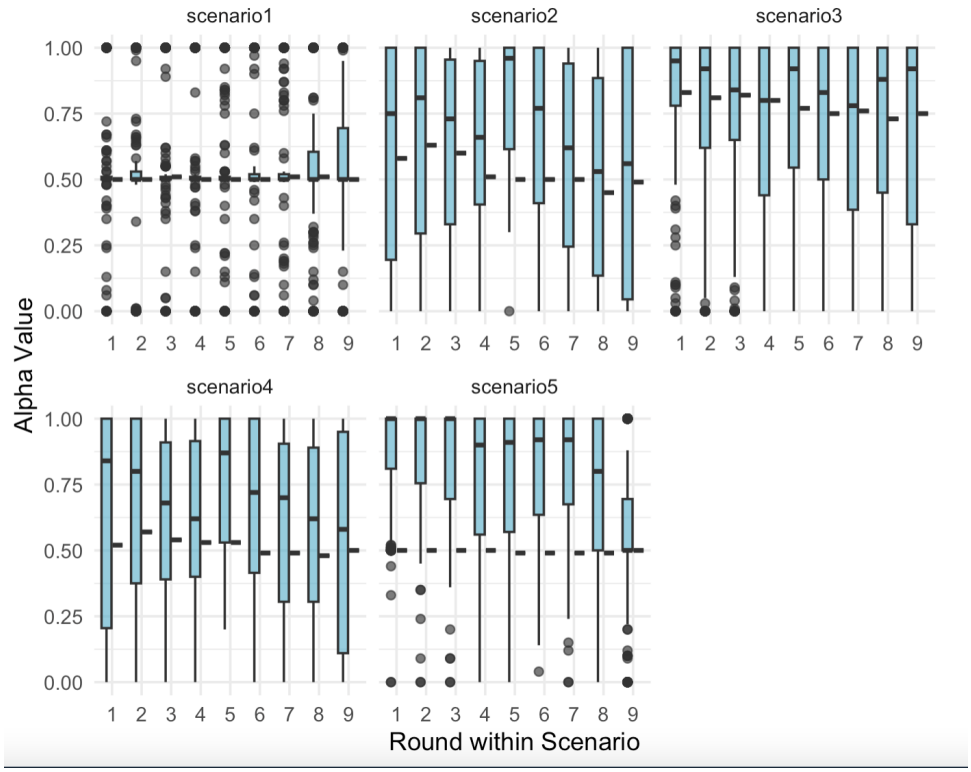
In three scenarios (2, 4, and 5), LPM_2 outperforms MV by clear margins on both MSE and MAE. These scenarios are characterized by asymmetric market conditions or consistently high α values (often ≈ 0.8), suggesting stronger behavioral deviations from the symmetric baseline. LPM_2's greater flexibility in capturing asymmetric preferences and tail-risk sensitivity appears to drive these gains.

Table 4: Scenario-wise performance

	Squared error			Absolute error			
Scenario	MV MSE	LPM_2 MSE	ΔMSE	MV MAE	LPM_2 MAE	ΔMAE	Winner
LPM_2 preferred							
scenario2	0.030	0.019	**-0.011**	0.171	0.127	**-0.044**	LPM_2
scenario4	0.017	0.010	**-0.007**	0.127	0.088	**-0.039**	LPM_2
scenario5	0.035	0.022	**-0.014**	0.181	0.141	**-0.040**	LPM_2
Tie							
scenario1	0.025	0.025	0.000	0.153	0.153	0.000	tie
MV preferred							
scenario3	0.019	0.046	**+0.027**	0.131	0.203	**+0.071**	MV

Scenario 1 (symmetric market) produces identical error metrics for both models (MSE = 0.025; MAE = 0.153). Here, participant α values are tightly centered near 0.5, both models predict 0.5, and variance is low—conditions under which MV suffices.

Figure 1: Distribution of Participant Choices For Round/Scenarios



Scenario 3 is the only case in which MV outperforms LPM_2, with MSE and MAE gaps of +0.027 and +0.071, respectively. By design, MV predicts allocations near 0.70, whereas LPM_2 produces higher targets around 0.90, gradually declining over rounds. Participant means align more closely with the MV level, contributing to its lower error. Notably, this scenario exhibits substantial "heaping" at the 0.5 allocation: 17.9% of observations fall within ±0.03 of 0.5, compared to only 2.1% in adjacent windows ($\chi^2 = 183.6$, $p < 0.001$).

This pattern is consistent with a half-and-half heuristic, suggesting that under these market conditions some participants may default to a salient focal point rather than optimally adjusting to asymmetries.

5.3 Win–Rate Analysis (Participant–Level Paired Losses)

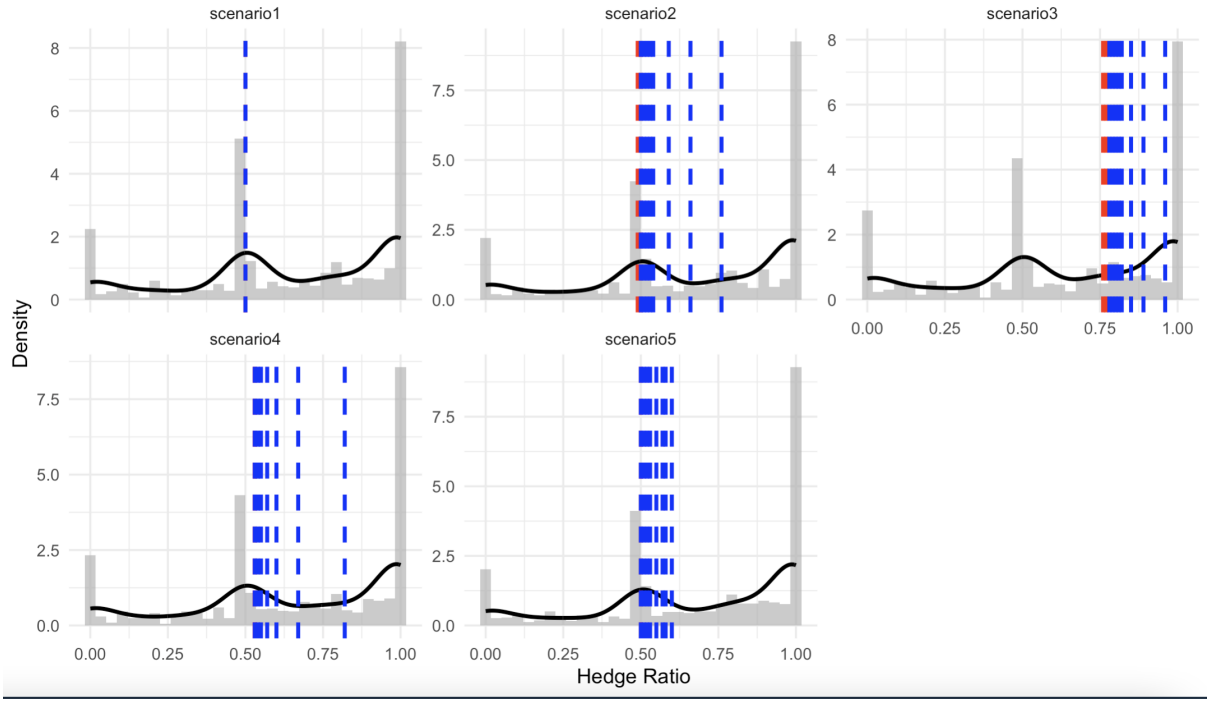
To align the unit of evaluation with individual decisions, we form a paired difference for every participant–round observation,

$$\Delta = \left| y - \hat{y}^{\mathrm{LPM}_2} \right| - \left| y - \hat{y}^{\mathrm{MV}} \right|,$$

where $\Delta < 0$ denotes a "win" for LPM_2 (closer to the observed choice y) and $\Delta > 0$ denotes a win for MV. We summarize results at two levels: (1) Decision-level win rates, counting each participant–round observation separately, and (2) Participant-level win rates, where each participant casts one "vote" per scenario based on their average absolute error over its nine rounds.

Figure 2 plots the distribution of participant allocations against model predictions for each scenario. In Scenarios 2, 4, and 5, the mass of choices lies closer to the higher LPM_2 targets, while in Scenario 3 the distribution shifts toward MV's prediction level, producing many small MV wins. The notable heaping at 0.5 in Scenario 3—17.9% of observations within ±0.03 versus 2.1% in adjacent bins ($\chi^2 = 183.6$, $p < 0.001$)—is consistent with a half-and-half heuristic, further explaining MV's advantage there.

Figure 2: Distribution of Participant Choices vs. Model Predictions by Scenario



Decision-level results (Table 5) confirm the aggregate error metrics. LPM_2 wins the majority of non-tied decisions in Scenarios 2 (54.3%), 4 (50.2%), and 5 (51.1%), while MV wins decisively in Scenario 3 (59.7%). Scenario 1 yields only ties because both models

Table 5: Win rates by scenario: decision-level and participant-level comparisons

Scenario	N	LPM_2 (%)	MV (%)	Tie (%)
Panel A: Decision-level wins				
scenario1	1332	0.0	0.0	100.0
scenario2	1332	54.3	34.4	11.3
scenario3	1332	40.2	59.7	0.1
scenario4	1332	50.2	38.7	11.1
scenario5	1332	51.1	26.4	22.4
Overall	**6660**	**39.2**	**31.8**	**29.0**
Panel B: Participant-level wins				
scenario1	148	0.0	0.0	100.0
scenario2	148	58.1	41.9	0.0
scenario3	148	32.4	66.9	0.7
scenario4	148	58.8	41.2	0.0
scenario5	148	70.3	27.7	2.0
Overall	**148**	**59.5**	**39.9**	**0.68**

Notes: Panel A reports the share of decisions in which each model is closer to the observed choice. Panel B reports the share of participants preferring each model based on average absolute error per scenario. "N" is the number of decisions in Panel A and the number of participants in Panel B.

predict identically. Pooled across all scenarios, LPM_2 wins 55.2% of non-tied decisions ($n = 4{,}729$, binomial $p = 1.22 \times 10^{-12}$).

Participant-level results (Table 5) tell the same story while down-weighting repeated measures: LPM_2 is preferred by 58.1% of participants in Scenario 2, 58.8% in Scenario 4, and 70.3% in Scenario 5; MV is preferred by 66.9% in Scenario 3; Scenario 1 is an all-tie case. Overall, LPM_2 is preferred by 59.5% of participants (55.3% after excluding ties).

These patterns underscore that while LPM_2 has a clear edge in Scenarios 2, 4, and 5, MV remains competitive—and in Scenario 3, superior—when participant allocations align more closely with its fixed prediction points. The 50–50 overall split in individual winners reflects substantial heterogeneity in preferences, but the scenario-level alignment with aggregate MSE/MAE results supports LPM_2 as the primary specification for this setting.

5.4 Individual Risk Perception

We estimate individual-level behavioral parameters under Cumulative Prospect Theory (CPT) using a regularized maximum-likelihood framework with the Prelec probability-weighting function [Pre98]. This approach recovers curvature over gains and losses, loss aversion, probability weighting, and choice noise from binary experimental choices, while stabilizing estimates through weakly informative Gaussian priors centered at canonical CPT values ($\alpha, \beta \approx 0.88$, $\lambda \approx 2.25$, $\rho \approx 0.69$).

5.4.1 CPT Parameter Estimation

For a mixed prospect with gain $d_1 \geq 0$ (probability q_1) and loss $d_2 \leq 0$ (probability q_2), the CPT value is

$$V(d_1, q_1; d_2, q_2 \mid \theta) = w^+(q_1)v(d_1) - w^-(q_2)\lambda v(-d_2), \tag{2}$$

where $\theta = (\alpha, \beta, \lambda, \rho, \sigma)$ collects the participant's parameters, $v(x) = x^\alpha$ for $x \geq 0$ and $v(x) = (-x)^\beta$ for $x < 0$, $\lambda > 0$ scales losses relative to gains, and $w^+(p)$, $w^-(p)$ are Prelec weights

$$w^\pm(p) = \exp\Big\{-\big(-\ln p\big)^\rho\Big\}, \qquad p \in (0,1). \tag{3}$$

The curvature parameters satisfy $\alpha, \beta \in (0,1]$ and the noise parameter $\sigma > 0$ governs choice stochasticity. Given two options A and B on trial t, with CPT values V_{it}^A and V_{it}^B computed from (2)–(3), we map values to choice probabilities using a logistic link with scale σ:

$$\Pr(Y_{it} = 1 \mid \theta_i) = \text{logit}^{-1}\left(\frac{V_{it}^A - V_{it}^B}{\sigma}\right) = \frac{1}{1 + \exp\big(-(V_{it}^A - V_{it}^B)/\sigma\big)}, \tag{4}$$

where $Y_{it} = 1$ denotes a choice of A.

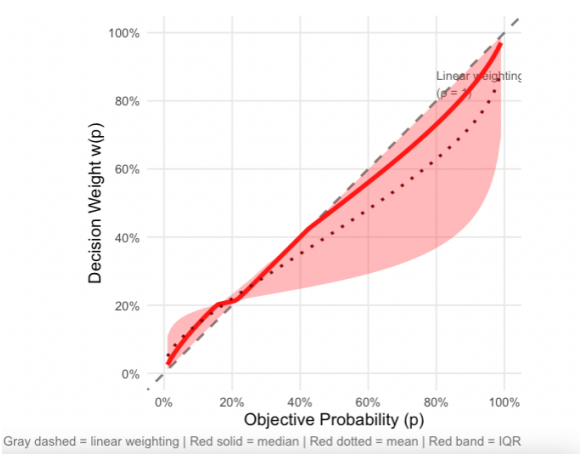
Parameters are identified for each participant by maximizing a *regularized* log-likelihood.

Figure 3 compares CPT and Constant Relative Risk Aversion (CRRA) fits using participant-level RMSE. Each point represents one participant: the x-axis is CRRA RMSE, the y-axis is CPT RMSE. Points below the diagonal indicate a CPT advantage.

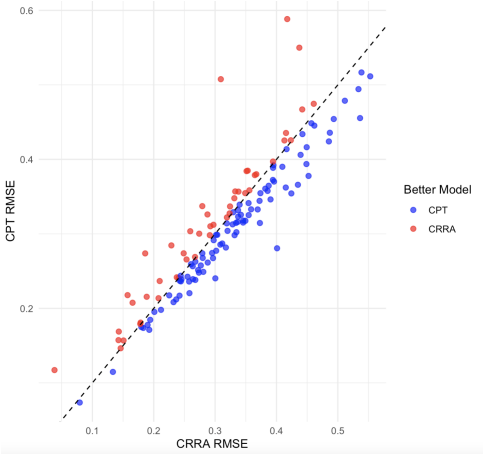
CPT outperforms CRRA for 93 of 140 participants (66.4%) and in 26 of 45 scenarios (57.8%). Most participants cluster between RMSE values of 0.15–0.45, with CPT generally lower. The positive correlation between model errors suggests some participants are inherently harder to fit, regardless of specification. CRRA occasionally performs better (red points above the diagonal), indicating that for a minority of participants, simpler curvature-based preferences capture behavior more effectively.

Overall, these results support CPT as the preferred modeling framework for hedging decisions, particularly because it accommodates loss aversion and probability weighting, while also revealing substantial heterogeneity in individual risk perception.

Figure 3: Probability weighting and model performance: CPT vs. EU



(a) Probability weighting



(b) Model performance by participant

6 Conclusion

This paper addresses a fundamental question in risk management: Do decision-makers hedge to minimize variance, or do they hedge to minimize downside risk? Our experimental evidence reveals that a substantial proportion of participants' hedging decisions align more closely with Lower Partial Moment (LPM_2) optimization than with the conventional Minimum Variance (MV) approach. This finding challenges the widespread use of variance-based measures in hedging models and supports incorporating downside risk measures that better reflect how individuals actually perceive and respond to risk.

The superior performance of LPM_2 stems from its exclusive focus on downside deviations—the outcomes that fall short of expectations—while ignoring upside fluctuations that are generally welcome. In contrast, MV's symmetric treatment of risk, which penalizes both gains and losses equally, systematically underestimates participants' loss aversion, particularly in scenarios with asymmetric risk profiles. Across most experimental conditions, LPM_2 provides more accurate predictions of observed hedge ratios, with its advantage most pronounced when downside protection is a primary concern.

Our findings also connect to the behavioral foundations underlying these risk measures. The preference for LPM_2-based hedging aligns with key predictions from cumulative prospect theory (CPT), particularly loss aversion and probability weighting. Participants systematically place greater weight on avoiding losses below their reference point than on achieving equivalent gains, consistent with the asymmetric value function in CPT. Moreover, the tendency to overweight small-probability extreme outcomes—a central feature of probability weighting—may explain why participants favor hedge ratios that provide stronger protection against tail risks. These behavioral factors suggest that LPM_2's descriptive success reflects not just a mathematical property but fundamental psychological tendencies in how individuals process risk.

These findings have important implications for hedging theory and practice. They suggest that traditional variance minimization may lead to systematically biased hedge ratio recommendations—either over-hedging and reducing beneficial upside volatility, or under-hedging and failing to adequately protect against extreme losses. More accurate models that incorporate asymmetric risk preferences could inform the design of hedging strategies that better align with decision-makers' actual risk perceptions, potentially improving both hedging effectiveness and adoption.

References

[Ede79] Ederington, L. H. "The Hedging Performance of the New Futures Markets". In: *Journal of Finance* 34 (1979), pp. 157–170.

[Joh60] Johnson, L. L. "The theory of hedging and speculation in commodity futures". In: *The Review of Economic Studies* 27.3 (1960), pp. 139–151.

[LVP17] Liu, P., Vedenov, D., and Power, G. J. "Is Hedging the Crack Spread No Longer All It's Cracked Up to Be?" In: *Energy Economics* 63 (2017), pp. 31–40. DOI: https://doi.org/10.1016/j.eneco.2017.01.020.

[Mar52] Markowitz, H. "Portfolio Selection". In: *The Journal of Finance* 7 (1) (1952), pp. 77–91. DOI: https://doi.org/10.2307/2975974.

[MGN08] Mattos, F., Garcia, P., and Nelson, C. “Relaxing standard hedging assumptions in the presence of downside risk”. In: *The Quarterly Review of Economics and Finance* 48 (1) (2008), pp. 78–93. DOI: https://doi.org/10.1016/j.qref.2006.01.003.

[MGP08] Mattos, F., Garcia, P., and Pennings, J. M. “Probability weighting and loss aversion in futures hedging”. In: *Journal of Financial Markets* 11 (4) (2008), pp. 433–452. DOI: https://doi.org/10.1016/j.finmar.2008.04.002.

[Pre98] Prelec, D. “The Probability Weighting Function”. In: *Econometrica* 66 (3) (1998), pp. 497–527. DOI: https://doi.org/10.2307/2998573.

[PV10] Power, G. J. and Vedenov, D. “Dealing with downside risk in a multi-commodity setting: A case for a “Texas hedge”?” In: *Journal of Futures Markets: Futures, Options, and Other Derivative Products* 30.3 (2010), pp. 290–304. DOI: https://doi.org/10.1002/fut.20411.

[TK79] Tversky, A. and Kahneman, D. “Prospect Theory : An Analysis of Decisions under Risk”. In: *Econometrica* 47 (1979), p. 278.

[TK92] Tversky, A. and Kahneman, D. “Advances in prospect theory: Cumulative representation of uncertainty”. In: *Journal of Risk and uncertainty* 5.4 (1992), pp. 297–323. DOI: https://doi.org/10.1007/BF00122574.

Annex: Scenario Design and Materials

This annex documents the five experimental scenarios. For each scenario, the table lists the round-by-round lottery for options L1 and L2, showing probabilities p and $1-p$ and the corresponding payoffs (low, high). Figures provide two panels: Panel A plots the payoff structure; Panel B visualizes the empirical distribution of hedge ratios and model benchmarks. All numbers are as implemented in the experiment.

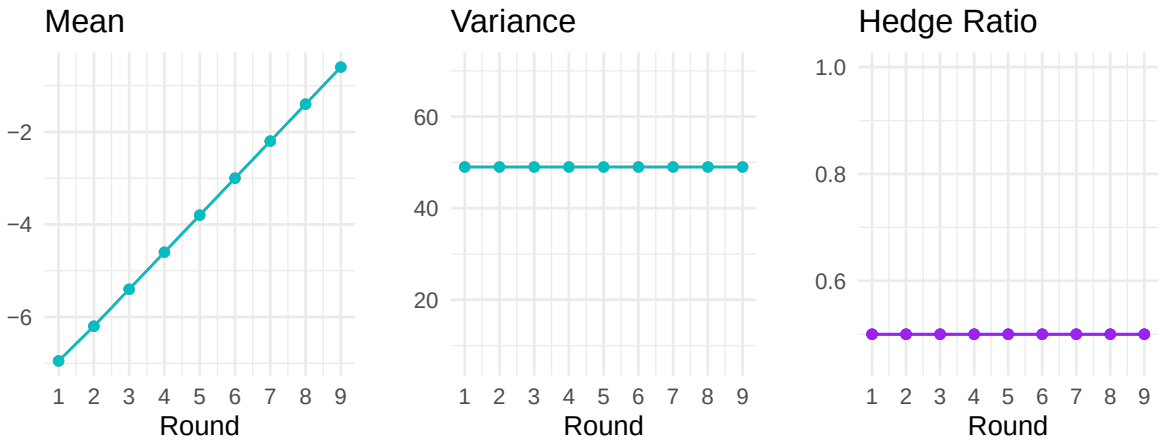
Scenario 1 (Symmetric benchmark)

The market is symmetric across options: probabilities and payoffs are matched for L1 and L2 in every round. Round means for L1 and L2 increase in parallel, and both MV and LPM_2 recommend a hedge ratio of 0.5 throughout.

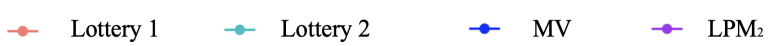
Table A1: Scenario 1: lottery by round

Round	$1-p$	$L1_{\text{high}}$	p	$L1_{\text{low}}$	p	$L2_{\text{high}}$	$1-p$	$L2_{\text{low}}$
round1	0.50	0.05	0.50	-13.95	0.50	0.05	0.50	-13.95
round2	0.50	0.80	0.50	-13.20	0.50	0.80	0.50	-13.20
round3	0.50	1.60	0.50	-12.40	0.50	1.60	0.50	-12.40
round4	0.50	2.40	0.50	-11.60	0.50	2.40	0.50	-11.60
round5	0.50	3.20	0.50	-10.80	0.50	3.20	0.50	-10.80
round6	0.50	4.00	0.50	-10.00	0.50	4.00	0.50	-10.00
round7	0.50	4.80	0.50	-9.20	0.50	4.80	0.50	-9.20
round8	0.50	5.60	0.50	-8.40	0.50	5.60	0.50	-8.40
round9	0.50	6.40	0.50	-7.60	0.50	6.40	0.50	-7.60

Figure A1: Scenario 1: design (Panel A) and behavior (Panel B)



Panel A: Payoff structure and probabilities.



Panel B: Distribution of hedge ratios vs. model predictions.

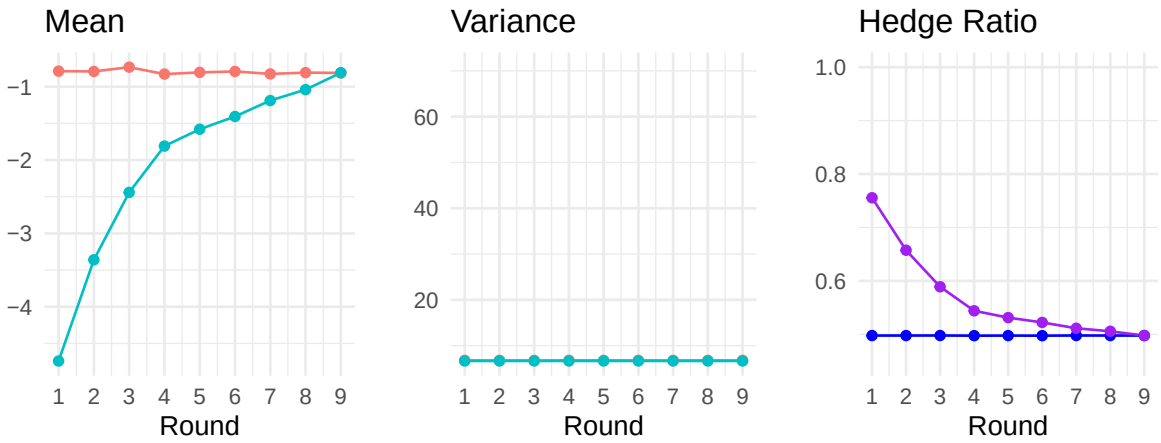
Scenario 2 (Asymmetric, equal variance)

The two options differ in probability–payoff profiles that evolve in opposite directions over rounds, while variances are matched. MV remains close to an equal allocation, whereas LPM_2 declines across rounds, reflecting sensitivity to downside risk asymmetry.

Table A2: Scenario 2: lottery by round

Round	$1-p$	$L1_{\text{high}}$	p	$L1_{\text{low}}$	p	$L2_{\text{high}}$	$1-p$	$L2_{\text{low}}$
round1	0.90	0.08	0.10	-8.60	0.10	3.00	0.90	-5.60
round2	0.80	0.51	0.20	-6.00	0.20	1.80	0.80	-4.65
round3	0.70	0.97	0.30	-4.71	0.30	1.50	0.70	-4.13
round4	0.60	1.30	0.40	-4.02	0.40	1.35	0.60	-3.92
round5	0.50	1.80	0.50	-3.41	0.50	1.00	0.50	-4.16
round6	0.40	2.40	0.60	-2.92	0.60	0.70	0.40	-4.57
round7	0.30	3.15	0.70	-2.53	0.70	0.50	0.30	-5.13
round8	0.20	4.40	0.80	-2.11	0.80	0.25	0.20	-6.20
round9	0.10	7.00	0.90	-1.68	0.90	0.05	0.10	-8.55

Figure A2: Scenario 2: design (Panel A) and behavior (Panel B)



Panel A: Payoff structure and probabilities.



Panel B: Distribution of hedge ratios vs. model predictions.

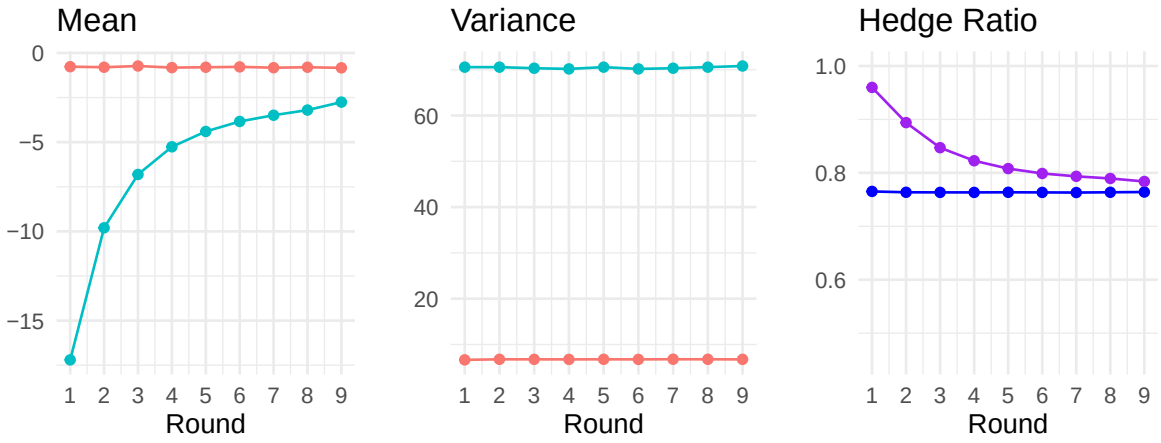
Scenario 3 (Higher mean L1, heavy left tail in L2)

Option L1 has a higher mean with moderate variance; L2 features more extreme negative outcomes and higher volatility. MV is constant by construction; LPM_2 shifts weight toward L1 and, in later rounds, moves closer to MV as L2's downside intensifies.

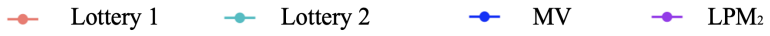
Table A3: Scenario 3: lottery by round

Round	$1-p$	$L1_{\text{high}}$	p	$L1_{\text{low}}$	p	$L2_{\text{high}}$	$1-p$	$L2_{\text{low}}$
round1	0.90	0.09	0.10	-8.50	0.10	8.00	0.90	-20.00
round2	0.80	0.50	0.20	-6.00	0.20	7.00	0.80	-14.00
round3	0.70	0.97	0.30	-4.70	0.30	6.00	0.70	-12.30
round4	0.60	1.30	0.40	-4.00	0.40	5.00	0.60	-12.10
round5	0.50	1.80	0.50	-3.40	0.50	4.00	0.50	-12.80
round6	0.40	2.40	0.60	-2.90	0.60	3.00	0.40	-14.10
round7	0.30	3.15	0.70	-2.53	0.70	2.00	0.30	-16.30
round8	0.20	4.40	0.80	-2.10	0.80	1.00	0.20	-20.00
round9	0.10	6.96	0.90	-1.70	0.90	0.05	0.10	-28.00

Figure A3: Scenario 3: design (Panel A) and behavior (Panel B)



Panel A: Payoff structure and probabilities.



Panel B: Distribution of hedge ratios vs. model predictions.

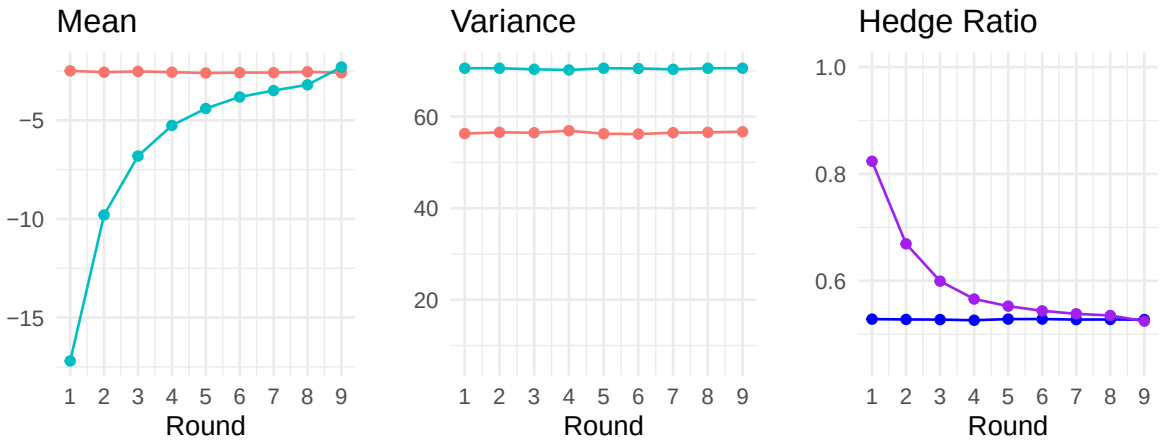
Scenario 4 (High variance, heavier losses in L2)

L1 ramps from very small to very large positive payoffs; L2 carries heavier left-tail risk, including extreme negatives. Early rounds imply stronger hedging against downside. MV remains stable; LPM_2 captures the asymmetric aversion to large losses.

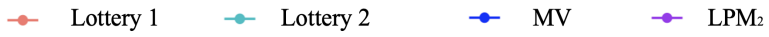
Table A4: Scenario 4: lottery by round

Round	$1-p$	$L1_{\text{high}}$	p	$L1_{\text{low}}$	p	$L2_{\text{high}}$	$1-p$	$L2_{\text{low}}$
round1	0.90	0.01	0.10	-25.00	0.10	8.00	0.90	-20.00
round2	0.80	1.20	0.20	-17.60	0.20	7.00	0.80	-14.00
round3	0.70	2.40	0.30	-14.00	0.30	6.00	0.70	-12.30
round4	0.60	3.60	0.40	-11.80	0.40	5.00	0.60	-12.10
round5	0.50	4.90	0.50	-10.10	0.50	4.00	0.50	-12.80
round6	0.40	6.60	0.60	-8.70	0.60	3.04	0.40	-14.10
round7	0.30	8.90	0.70	-7.50	0.70	2.00	0.30	-16.30
round8	0.20	12.50	0.80	-6.30	0.80	1.00	0.20	-20.00
round9	0.10	20.00	0.90	-5.00	0.90	0.50	0.10	-27.50

Figure A4: Scenario 4: design (Panel A) and behavior (Panel B)



Panel A: Payoff structure and probabilities.



Panel B: Distribution of hedge ratios vs. model predictions.

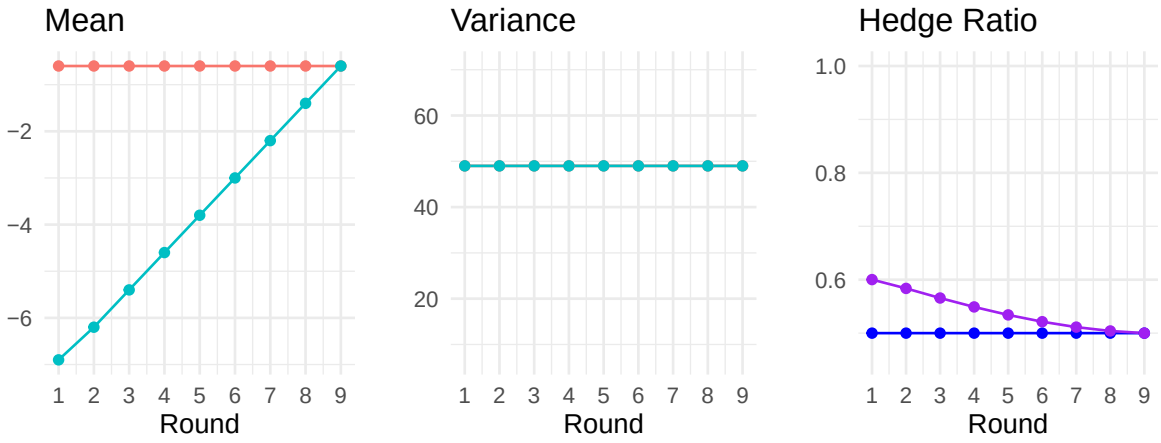
Scenario 5 (Constant $p = 0.5$, rising L2 upside)

Probabilities remain fixed at 0.5, allowing a focused comparison of how MV and LPM_2 respond to changing payoff magnitudes. MV stays at 0.5 by construction; LPM_2 declines as L2's profile evolves, consistent with greater attention to downside risk.

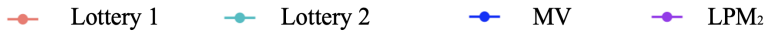
Table A5: Scenario 5: lottery by round

Round	$1-p$	$L1_{\text{high}}$	p	$L1_{\text{low}}$	p	$L2_{\text{high}}$	$1-p$	$L2_{\text{low}}$
round1	0.50	6.40	0.50	-7.60	0.50	0.10	0.50	-13.90
round2	0.50	6.40	0.50	-7.60	0.50	0.80	0.50	-13.20
round3	0.50	6.40	0.50	-7.60	0.50	1.60	0.50	-12.40
round4	0.50	6.40	0.50	-7.60	0.50	2.40	0.50	-11.60
round5	0.50	6.40	0.50	-7.60	0.50	3.20	0.50	-10.80
round6	0.50	6.40	0.50	-7.60	0.50	4.00	0.50	-10.00
round7	0.50	6.40	0.50	-7.60	0.50	4.80	0.50	-9.20
round8	0.50	6.40	0.50	-7.60	0.50	5.60	0.50	-8.40
round9	0.50	6.40	0.50	-7.60	0.50	6.40	0.50	-7.60

Figure A5: Scenario 5: design (Panel A) and behavior (Panel B)



Panel A: Payoff structure and probabilities.



Panel B: Distribution of hedge ratios vs. model predictions.

Experiment Screen

Figure A6: Experimental Decision Screen Shown to Participants

