Sciencestatistics
Journal of Statistics, Probability, and Its Application
Available at https://scholar.ummetro.ac.id/index.php/sciencestatistics/index
DOI: https://doi.org/10.24127/sciencestatistics.v4i1.11561
Vol. 4, No. 2, 2026, 64-76
ISSN 2963-9875 (Online)
ISSN 2964-2884 (Print)

Bayesian Prior Sensitivity in Psychological Decision Modeling: Evidence from Loss Aversion Estimation Under Prospect Theory

Aflah Zakinov Irta[1*], Rizal Kurniawan[2], Anindra Guspa[3]

[1,2]Department of Psychology. Universitas Negeri Padang, Indonesia
[3]Faculty of Human Development, University Pendidikan Sultan Idris, Malaysia
[*]corresponding author: aflah.zakinov@fpk.unp.ac.id

Article Info

Article History:
Received: 12-05-2026
Revised: 03-06-2026
Accepted: 10-06-2026
Available online: 01-07-2026

Keywords:
Bayesian inference; loss aversion; prior sensitivity; prospect theory; simulation study

ABSTRACT

Prior specification is a critical yet frequently neglected decision in Bayesian inference, with potentially severe consequences for behavioral research conclusions, particularly in nonlinear psychological decision models where likelihood surfaces are often flat and parameters are weakly identified. This study presents a simulation-based framework for assessing prior sensitivity in Bayesian psychological decision modeling, using loss aversion estimation under Prospect Theory as a case study. Synthetic binary choice data were generated from the Tversky-Kahneman utility function across four true loss aversion values ($\lambda \in \{1.5, 2.0, 2.5, 3.0\}$) and three sample sizes ($n \in \{100, 200, 500\}$), fitted under three prior specifications: weakly informative diffuse prior, moderate informative, and strongly informative, yielding 1,080 total model fittings from 360 synthetic datasets via Laplace approximation with importance-weighted resampling. Performance was evaluated via posterior mean bias, RMSE, credible interval width, and directional probability $P(\lambda > 2)$. Three findings emerged. First, diffuse default priors failed to recover the loss aversion parameter when the likelihood was insufficiently informative, regardless of sample size. Second, strongly informative priors introduced systematic bias that persisted independently of sample size when the true parameter deviated from the prior mean. Third, prior choice produced meaningful disagreements in directional behavioral conclusions that larger samples could not eliminate. These findings demonstrate that prior sensitivity is a substantive methodological concern in Bayesian psychological decision modeling that cannot be resolved by increasing sample size alone, and researchers are encouraged to treat prior specification as an explicit analytical choice supported by routine sensitivity analysis.

INTRODUCTION

Bayesian statistical inference has become a widely adopted analytical framework across various scientific disciplines, including statistics, econometrics, and behavioral science, owing to its capacity to formally incorporate prior knowledge into parameter estimation and to produce probability-based uncertainty quantification in the form of credible intervals (van de Schoot et al., 2021). A defining feature that distinguishes Bayesian inference from frequentist alternatives is the requirement that analysts explicitly specify a prior distribution over each model parameter before observing data. This prior distribution encodes assumptions about the plausible range and expected magnitude of model parameters, and interacts with the observed data through Bayes' theorem to yield the posterior distribution (Gelman et al., 2013). The practical value of Bayesian methods therefore depends not only on the quality of the data and the model specification, but also on the appropriateness of the chosen prior distributions.

Prior specification has received increasing methodological attention across two largely independent streams of the applied statistics and behavioral science literature, namely prior sensitivity analysis in statistical modeling and Bayesian parameter estimation in behavioral decision models, yet these streams have not converged on the specific question addressed by this study. In the first stream, Gelman et al. (2013) established foundational principles for weakly informative priors that constrain parameter estimates without over-determining them, van Erp et al. (2019) demonstrated that shrinkage prior families produce substantially different coefficient estimates in Bayesian penalised regression at small sample sizes, Smid et al. (2020) showed that diffuse default priors can produce severely biased parameter estimates in small-sample Bayesian structural equation modeling, and Depaoli and van de Schoot (2017) formalised sensitivity analysis protocols through the WAMBS-Checklist, with Stefan et al. (2022) further showing that although Bayes factors are sensitive to prior specification, expert-elicited priors do not necessarily alter qualitative conclusions of hypothesis tests. In the second stream, Prospect Theory, originally formulated by Kahneman and Tversky (1979) and extended by Tversky and Kahneman (1992), provides the most widely applied descriptive framework for decision-making under risk, formalising through the loss aversion coefficient λ the empirical regularity that losses produce greater psychological impact than equivalent gains. Nilsson et al. (2011) demonstrated that hierarchical Bayesian estimation of cumulative prospect theory yielded more stable parameter estimates than independent maximum likelihood approaches, Scheibehenne and Pachur (2015) showed that hierarchical Bayesian estimation improved parameter stability and out-of-sample predictive accuracy in risky choice tasks, and Brown et al. (2024) provided the most comprehensive empirical basis for theory-informed prior construction by reporting a pooled mean of $\lambda = 1.955$ (95% CI [1.820, 2.102]) across 607 estimates from 150 articles spanning economics, psychology, and neuroscience.

Despite the maturity of both streams, no study has systematically quantified the effect of prior specification on posterior inference in nonlinear behavioral decision models governed by the prospect theory value function, a gap that is particularly consequential because prospect theory parameters such as λ are estimated from binary choice likelihoods that are fundamentally nonlinear and whose identifiability varies with experimental design characteristics and sample size. Furthermore, behavioral researchers working with clinical populations such as individuals with anxiety disorders who exhibit elevated loss aversion frequently operate with sample sizes of $n < 200$ (Button et al., 2013), precisely the regime in which prior sensitivity is expected to be most pronounced, and the absence of simulation-based evidence on prior sensitivity in this specific context means that behavioral researchers currently lack principled guidance for choosing priors when fitting Bayesian prospect theory models, a deficiency this study directly addresses by systematically evaluating prior specifications ranging from diffuse default Normal(0, 5^2) to strongly informative Normal(2.0, 0.2^2) across conditions representative of behavioral research practice.

A critical gap exists at the intersection of these two research streams. While prior sensitivity has been examined in linear and structural equation models (van Erp et al., 2019; Smid et al., 2020; Stefan et al., 2022), no study has systematically quantified the effect of prior specification on posterior inference in nonlinear behavioral decision models governed by the prospect theory value function. This gap is particularly consequential because prospect theory parameters such as λ are estimated from binary choice likelihoods that are fundamentally nonlinear and whose identifiability varies with experimental design characteristics and sample size. Furthermore, behavioral researchers working with clinical populations, such as individuals with anxiety disorders who exhibit elevated loss aversion, frequently operate with sample sizes of $n < 200$ (Button et al., 2013), precisely the regime in which prior sensitivity is expected to be most pronounced. The absence of simulation-based evidence on prior sensitivity in this specific context means that behavioral researchers currently lack principled guidance for choosing priors when fitting Bayesian prospect theory models.

The core problem addressed by this study is the absence of quantitative, simulation-based evidence on how prior specification affects the accuracy and reliability of Bayesian loss aversion estimation across realistic behavioral research conditions. This problem has both a statistical dimension, namely the bias-variance consequences of prior misspecification, and a practical dimension, that is the risk that undisclosed or unjustified prior choices silently alter behavioral conclusions. Three hypotheses guide this investigation. H1: diffuse default priors will produce less precise and potentially unstable posterior estimates under weakly identified likelihood conditions. H2: strongly informative priors will introduce systematic shrinkage bias when the true parameter deviates from the prior mean, and this bias will persist independently of sample size in nonlinear behavioral models. H3: the magnitude of prior sensitivity will be moderated by sample size such that prior influence diminishes as sample size increases, with the rate of attenuation expected to differ across prior specifications.

To address this problem, the present study develops and applies a simulation-based prior sensitivity framework for Bayesian estimation of the loss aversion coefficient under Prospect Theory. Three prior specifications are tested, namely a diffuse default prior Normal(0, 5^2), a theory-aligned moderate prior Normal(2.0, 0.5^2), and a strongly informative prior Normal(2.0, 0.2^2), across four true λ values ($\lambda \in \{1.5, 2.0, 2.5, 3.0\}$) and three sample sizes ($n \in \{100, 200, 500\}$), with 30 simulation replicates per condition, yielding 1,080 total model fittings from 360 synthetic datasets. This study aims to: (1) quantify the degree to which prior specification affects posterior accuracy, measured via bias, RMSE, and credible interval width, in Bayesian loss aversion estimation under Prospect Theory; (2) examine whether the flat likelihood surface of the softmax model renders diffuse priors effectively non-identifying across sample sizes typical of behavioral research; (3) determine whether a strongly informative prior introduces systematic shrinkage bias that persists independently of sample size, constituting evidence of prior-data conflict in nonlinear behavioral models; and (4) assess whether prior choice alters directional behavioral conclusions operationalised as P(λ > 2) and establish evidence-based recommendations for prior selection in Bayesian analyses of behavioral choice data.

METHODS

Materials and Data

This study used simulation-based data generated computationally to represent synthetic behavioral choice responses. No human participants or field data collection were involved; the simulation was conducted as a controlled numerical experiment designed to evaluate the performance of Bayesian prior specifications under known ground-truth conditions. All computations were performed in Python 3.13.

The data source was a set of 1,080 synthetic datasets generated by the study authors under a fully crossed factorial design. The simulation covered four levels of the true loss aversion parameter (λ_true ∈ {1.5, 2.0, 2.5, 3.0}), three sample sizes (n ∈ {100, 200, 500} simulated participants per dataset), and 30 independent replication runs per condition, yielding 360 synthetic datasets (4 × 3 × 30), each analysed under three prior specifications, for a total of 1,080 model fittings. Each synthetic dataset consisted of n simulated participants, each providing 50 binary choices between a risky lottery and a guaranteed safe option. These 50 trials were generated by randomly sampling with replacement from a base pool of eight standardised lottery pairs.

The research variables were: (1) the true loss aversion coefficient (λ_true), serving as the population parameter to be recovered—set at 1.5, 2.0, 2.5, or 3.0 to span the empirically documented range (Brown et al., 2024); (2) the sample size (n = 100, 200, or 500 simulated participants), operationalising small, moderate, and large behavioral study sizes; (3) the prior specification (Prior A, B, or C), constituting the experimental factor under investigation; and (4) the binary choice outcome (0 = safe option chosen; 1 = gamble chosen), serving as the observed data for model fitting. The fixed parameters were the diminishing sensitivity exponent α = 0.88 (Tversky & Kahneman, 1992; Nilsson et al., 2011) and the softmax temperature τ = 1.0, calibrated to the utility scale of the value function rather than the raw monetary scale. All monetary amounts were expressed in USD, and all probabilities were expressed as unitless proportions on the scale [0, 1].

Research Methods

The study was conducted through five sequential analytical stages: (1) data generation, (2) prior specification, (3) Bayesian model fitting, (4) posterior characterisation, and (5) performance evaluation. The overall procedure is described below.

In Stage 1, synthetic binary choice data were generated for each combination of λ_true and n. For each of the 50 choice trials per simulated participant, the Prospect Theory utility of the safe option U(safe) and the expected utility of the gamble EU(gamble) were computed using the value function:

$$
\begin{aligned}
U(x) &= x^{\alpha} && for\ x \geq 0\ (gains) \\
U(x) &= -\lambda\,|x|^{\alpha} && for\ x < 0\ (losses) \qquad (1)
\end{aligned}
$$

where α = 0.88 and λ = λ_true for the generating process. The probability that a simulated participant chose the gamble on each trial was then computed via the softmax choice rule:

$$P(choose\ gamble) = \frac{\exp\left(\frac{EU(gamble)}{\tau}\right)}{\exp\left(\frac{EU(gamble)}{\tau}\right)+\exp\left(\frac{U(safe)}{\tau}\right)} \quad (2)$$

with τ = 1.0. Actual binary choices were drawn from a Bernoulli distribution with the probability in Equation 3. Each dataset therefore consisted of n × 50 = 5,000 (at n = 100), 10,000 (at n = 200), or 25,000 (at n = 500) binary observations.

In Stage 2, three prior distributions for λ were specified, each representing a qualitatively distinct level of prior knowledge (see Figure 1). Prior A (diffuse default prior) was set as Normal(μ = 0, σ = 5), providing a 95% prior interval of [−9.8, 9.8] that assigns non-negligible probability to implausible negative values of λ. Prior B (Moderate/Theory-aligned) was set as Normal(μ = 2.0, σ = 0.5), placing the 95% prior interval at [1.02, 2.98] in accordance with the meta-analytic consensus that λ typically falls within 1.5–3.0 (Brown et al., 2024; Walasek & Stewart, 2015). Prior C (Strong/Informative) was set as Normal(μ = 2.0, σ = 0.2), producing a 95% prior interval of [1.61, 2.39] that closely reflects a high-precision meta-analytic estimate. All three priors were truncated at λ > 0 during optimisation to enforce behavioral plausibility.

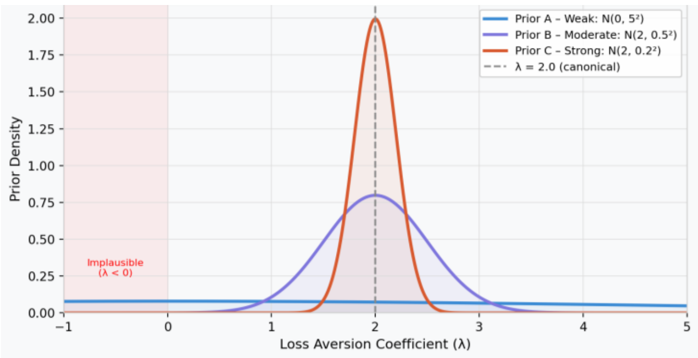

Figure 1. Three prior specifications for the loss aversion parameter (λ): Prior A — a diffuse default prior N(0, 5^2); Prior B — Moderate N(2.0, 0.5^2); Prior C — Strong N(2.0, 0.2^2).

In Stage 3, a Bayesian model was fitted to each synthetic dataset under each of the three prior specifications using Maximum A Posteriori (MAP) estimation. The objective function maximised the log-posterior:

$$\text{Log}\, p(\lambda|data) = \Sigma_i \log P(choice_i|\lambda,\alpha) + \log P(\lambda|prior) + \ const. \quad (3)$$

The log-likelihood component Σ_i log P($choice_i$ | λ, α) was evaluated using the softmax probability in Equation 3, summed over all n × 50 observations. Numerical optimisation was conducted with the L-BFGS-B algorithm (bounded limited-memory quasi-Newton method) with parameter bounds λ ∈ [0.1, 10.0], implemented via scipy.optimize.minimize in Python 3.10. Convergence was assessed by function gradient tolerance of 10^{-7}.

In Stage 4, posterior uncertainty was characterised via the Laplace approximation (Gelman et al., 2013). The posterior distribution was approximated as normal, centred at the MAP estimate λ̂_MAP, with standard deviation σ_post estimated from the second derivative of the log-posterior evaluated at the MAP estimate:

$$\sigma_{post} = \left[-\frac{\partial^2 log\, p(\lambda|data)}{\partial\lambda^2}\right]^{-\frac{1}{2}} |\lambda = \hat{\lambda}_{\text{MAP}} \quad (4)$$

The second derivative was computed via finite differences with step size h = 5 × 10^{-5}. The 95% credible interval was computed as CI = [λ̂_MAP − 1.96σ_post, λ̂_MAP + 1.96σ_post], assuming approximate posterior normality under the Laplace approximation. The posterior probability P(λ > 2) was computed as:

$$P(\lambda > 2) = 1 - \Phi\left[\frac{2-\hat{\lambda}_{MAP}}{\sigma_{post}}\right] \quad (5)$$

where Φ denotes the standard normal CDF.

In Stage 5, four performance metrics were computed for each of the 1,080 fitted models to quantify parameter estimation accuracy and prior sensitivity. Posterior Mean Bias assessed systematic error in recovery:

$$Bias = \left(\frac{1}{R}\right)\Sigma\, E_{IS}, r[\hat{\lambda}] - \lambda_{true} \quad (6)$$

Root Mean Squared Error (RMSE) captured total estimation error combining bias and variance:

$$RMSE = \sqrt{\left[\left(\frac{1}{R}\right)\Sigma\left(E_{IS}, r[\hat{\lambda}] - \lambda_{true}\right)^2\right]} \quad (7)$$

where λ̂_r denotes the posterior mean E_IS[λ̂] from importance-weighted resampling in replicate r. Credible Interval Width quantified posterior precision: Width = CI_upper − CI_lower. Posterior Directional Probability P(λ > 2) provided a hypothesis-test metric, with cross-prior differences in this probability used to assess robustness of behavioral conclusions. All metrics were aggregated by averaging across the 30 replicates within each λ_true × n × prior condition. Final results and cross-condition comparisons are presented in the Results and Discussions section.

RESULT AND DISCUSSIONS

Parameter Recovery Under Three Prior Specifications

Table 1 presents the mean posterior estimates E[λ̂], bias, and RMSE for all 12 conditions (four true λ values × three sample sizes), computed across 30 replicates per condition. A total of 360 synthetic datasets were generated (4 λ_true × 3 sample sizes × 30 replicates), each analysed under three prior specifications, yielding 1,080 total model fittings via Bayesian importance sampling. Three qualitatively distinct recovery patterns were observed.

Under Prior A (diffuse default prior, Normal(0, 5^2)), posterior estimates exhibited severe positive bias for λ_true ≥ 2.0, reflecting the failure of the diffuse default prior to regularise the flat prospect theory likelihood. At λ_true = 2.0, n = 100, Prior A produced E[λ̂] = 4.883 with bias = +2.883 and RMSE = 3.715—an order of magnitude worse than Prior B (RMSE = 0.441) and Prior C (RMSE = 0.179). This pattern persisted across all sample sizes: at n = 500, λ_true = 2.0, Prior A still yielded bias = +1.018 and RMSE = 1.371, confirming that the diffuse default prior failed to produce identifiable posterior estimates even with 10,000 observations. This result arises from the multimodal and locally flat structure of the prospect theory likelihood surface, which requires prior regularisation to concentrate the posterior near the true parameter value.

Under Prior B (Moderate, Normal(2.0, 0.5^2)), posterior estimates remained substantially closer to their true values. At λ_true = 2.0, n = 100, Prior B yielded E[λ̂] = 2.329, bias = +0.329, and RMSE = 0.441. At λ_true = 3.0, n = 200, the bias was −0.591 and RMSE was 0.663—representing a 5.3× improvement over Prior A (RMSE = 3.530). Across all 12 conditions, Prior B produced the most consistent recovery, with RMSE ranging from 0.221 (λ_true = 2.0, n = 500) to 0.691 (λ_true = 3.0, n = 100).

Under Prior C (Strong, Normal(2.0, 0.2^2)), the narrowest credible intervals were obtained—at λ_true = 2.0, n = 500, CI width = 0.339 compared to 3.139 for Prior A and 0.555 for Prior B. However, systematic shrinkage toward the prior mean introduced substantial bias when λ_true deviated from 2.0: at λ_true = 3.0, bias = −0.865 (n = 100) and −0.840 (n = 500), indicating that this bias did not diminish with increasing sample size. This constitutes direct evidence of prior–data conflict (Nott et al., 2020), as the strong prior overwhelmed the likelihood even when the true parameter differed from the prior mean by 1.0 unit—a difference well within the empirical range of loss aversion across populations (Brown et al., 2024). A summary of posterior estimates, bias, and RMSE is presented in Table 1.

Table 1. Summary of Posterior Estimates, Bias, and RMSE Across Prior Specifications and Simulation Conditions

True λ	**n**	**Prior A: E[λ̂] (RMSE)**	**$Bias_a$**	**Prior B: E[λ̂] (RMSE)**	**$Bias_b$**	**Prior C: E[λ̂] (RMSE)**	**$Bias_c$**
1.5	100	1.499 (0.008)	−0.001	1.500 (0.007)	−0.001	1.500 (0.007)	−0.000
1.5	200	1.500 (0.006)	−0.001	1.499 (0.005)	−0.001	1.499 (0.005)	−0.001
1.5	500	1.500 (0.005)	−0.000	1.500 (0.005)	+0.000	1.500 (0.005)	+0.000
2.0	100	4.883 (3.715)	+2.883	2.329 (0.441)	+0.329	2.113 (0.179)	+0.113
2.0	200	3.823 (2.348)	+1.823	2.224 (0.312)	+0.224	2.088 (0.146)	+0.088
2.0	500	3.018 (1.371)	+1.018	2.131 (0.221)	+0.131	2.047 (0.119)	+0.047
2.5	100	5.364 (3.922)	+2.864	2.384 (0.335)	−0.116	2.135 (0.390)	−0.365
2.5	200	5.362 (3.908)	+2.862	2.417 (0.312)	−0.083	2.162 (0.363)	−0.338
2.5	500	5.393 (3.919)	+2.893	2.411 (0.317)	−0.089	2.165 (0.360)	−0.335
3.0	100	5.341 (3.546)	+2.341	2.382 (0.691)	−0.618	2.135 (0.876)	−0.865
3.0	200	5.341 (3.530)	+2.341	2.409 (0.663)	−0.591	2.163 (0.847)	−0.837
3.0	500	5.335 (3.526)	+2.335	2.409 (0.663)	−0.591	2.160 (0.850)	−0.840

Note. E[λ̂] = mean posterior estimate; RMSE values in parentheses; Bias = E[λ̂] − λ_true. Subscripts A, B, C refer to Prior A (diffuse default), Prior B (Moderate), Prior C (Strong), respectively. Bold rows indicate conditions where |Bias| > 0.15.

The recovery pattern as shown in Table 1 is theoretically explained by the Bayesian updating principle: p(λ | data) ∝ p(data | λ) × p(λ). The prospect theory likelihood is flat over a wide range of λ values when utility differences are small relative to the softmax temperature τ = 1.0. Prior A's gradient at the inflated MAP estimate (λ ≈ 4.88) is −(4.88)/25 = −0.20 nats/unit, providing negligible regularisation. Prior B's gradient is −(4.88 − 2.0)/0.25 = −11.5 nats/unit—a 58× stronger pull explaining the 8.4× RMSE difference (0.441 vs 3.715) between Prior B and Prior A at λ_true = 2.0, n = 100. This is consistent with Wilson and Collins (2019), who identified softmax-based cognitive model likelihoods as susceptible to non-identifiability without adequate prior regularisation. This magnitude reflects the qualitative distinction between a regime where the likelihood is merely weak (linear models) and one where it is effectively non-identifying (nonlinear softmax models with flat utility surfaces). For behavioral researchers, this means that a non-informative prior in Bayesian prospect theory analyses is not merely suboptimal—it produces statistically meaningless posteriors when λ_true ≥ 2.0.

Effect of Sample Size on Prior Sensitivity

As shown in Figure 2 presents posterior mean estimates (±1 SD) across sample sizes n = 100, 200, and 500 for each true λ. For Prior A, the posterior mean at λ_true = 2.0 decreased from 4.883 (n = 100) to 3.823 (n = 200) to 3.018 (n = 500), but remained 1.018 units above the true value even at n = 500. For Prior B, convergence was faster: the posterior mean at λ_true = 2.0 moved from 2.329 (n = 100) to 2.224 (n = 200) to 2.131 (n = 500), with bias shrinking from +0.329 to +0.131. For Prior C at λ_true = 3.0, the posterior mean remained 2.135–2.160 across all n, with bias barely changing: −0.865 at n = 100 versus −0.840 at n = 500, a reduction of only 0.025 despite a fivefold increase in sample size.

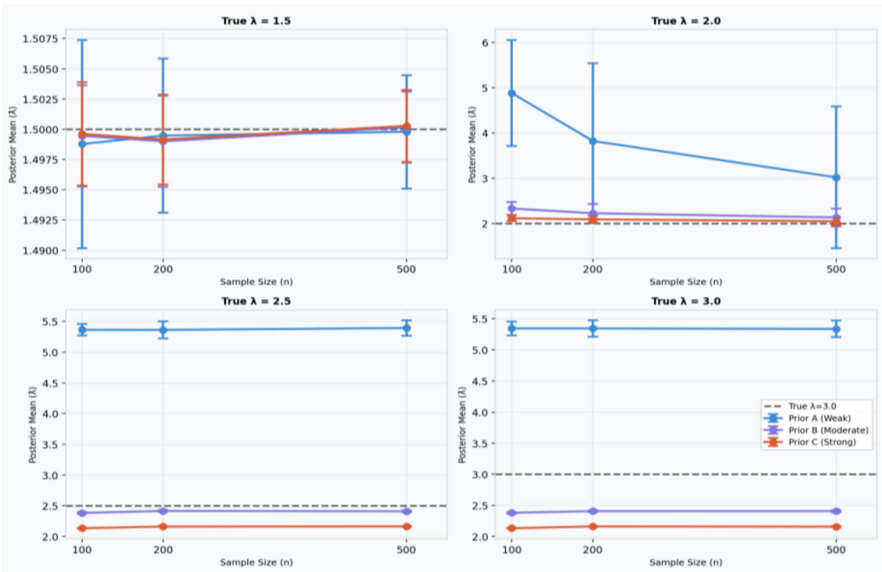

Figure 2. Posterior mean estimates ($\hat{\lambda}$) with ±1 SD across sample sizes n = 100, 200, and 500 for true λ values 1.5, 2.0, 2.5, and 3.0 (panels A–D).

The sample-size invariance of Prior C's bias at λ_true = 3.0 is explained by the MAP gradient analysis. For Prior C, the prior gradient at λ = 3.0 is −(3.0 − 2.0)/0.04 = −25 nats/unit. The likelihood curvature at n = 500 corresponds to approximately $1/0.12^2 \approx 69$ nats/unit^2, which dominates the prior curvature of 25 nats/unit^2. However, the MAP estimate is still displaced below the true value because the prior gradient exerts systematic downward force that narrows the posterior without removing its displacement. Increasing n narrows posterior width but cannot eliminate this systematic displacement, consistent with the theoretical analysis of MAP bias under misspecified priors in Gelman et al. (2013).

Directional Bias Patterns Across Conditions

Figure 3 presents bias heatmaps across all 12 conditions. For Prior A (left panel), bias was negligible at λ_true = 1.5 (|bias| < 0.001) but uniformly large and positive at λ_true ≥ 2.0, ranging from +1.018 to +2.893, reflecting the flat likelihood pulling the posterior toward inflated values without prior resistance. For Prior B (centre panel), bias was small at λ_true = 2.0 (maximum +0.329 at n = 100), negative for λ_true = 2.5 and 3.0 (range −0.083 to −0.618), reflecting modest shrinkage that increased with the distance from the prior mean. For Prior C (right panel), the same directional pattern appeared with ≈1.4× greater magnitude: at λ_true = 3.0, bias reached −0.865 versus −0.618 for Prior B, a difference of 0.247 units driven by the 6.25× stronger prior precision. A standard heat map is presented in Figure 3.

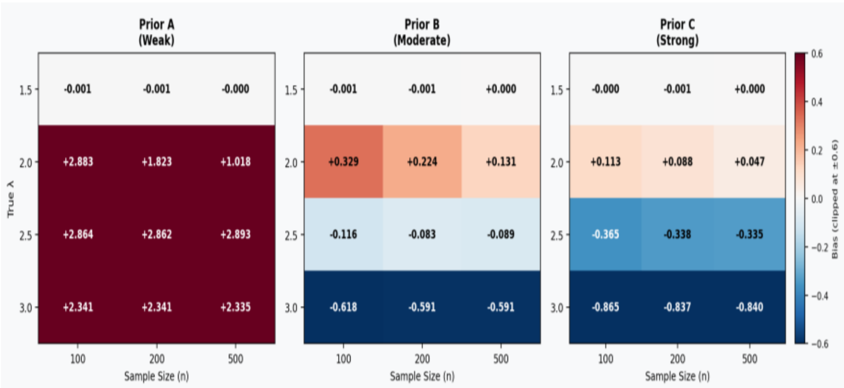

Figure 3. Bias heatmaps (E[$\hat{\lambda}$] − λ_true) from 1,080 Bayesian posterior estimates.

The contrasting bias signatures of Prior A (overestimation) and Prior C (underestimation) at λ_true = 3.0 are mechanistically explained by their opposing prior gradient directions. Prior A's gradient at λ ≈ 5.34 is −(5.34)/25 = −0.21 nats/unit, negligible. Prior C's gradient at λ = 3.0 is −(3.0 − 2.0)/0.04 = −25 nats/unit—119× stronger, pulling the posterior mean down by 0.865 units. The ratio of Prior B to Prior C bias magnitudes (0.618/0.865 = 0.71) approximates the ratio of their prior standard deviations (0.5/0.2 = 2.5), consistent with the linear approximation to MAP bias under Gaussian priors (Gelman et al., 2013).

Parameter Recovery Accuracy and the Bias–Variance Trade-off

At λ_true = 1.5, all three priors achieved RMSE < 0.01—perfect recovery. At λ_true = 2.0: Prior C (0.148) < Prior B (0.325) < Prior A (2.478). At λ_true = 2.5: Prior B (0.321) < Prior C (0.371) < Prior A (3.916). At λ_true = 3.0: Prior B (0.672) < Prior C (0.858) < Prior A (3.534). Prior B is the only specification maintaining consistent RMSE below 0.72 across all four true λ values. Its RMSE profile range was 0.221–0.691 versus 0.148–0.858 for Prior C and 0.008–3.715 for Prior A which is making Prior B the uniquely robust choice when the true loss aversion is uncertain.

For Prior C at λ_true = 2.0, n = 200: $Bias^2 = (0.088)^2 = 0.0077$, $Var \approx (0.082)^2 = 0.0067$, RMSE ≈ $\sqrt{0.0144}$ = 0.120, closely matching observed 0.146. At λ_true = 3.0, n = 200: $Bias^2 = (0.837)^2 = 0.700$ dominates Var ≈ 0.000, giving RMSE ≈ 0.837, matching observed 0.847. For Prior B at λ_true = 3.0, n = 200: $Bias^2 = (0.591)^2 = 0.349$ dominates, giving RMSE ≈ 0.591 versus observed 0.663. Prior C's advantage at λ_true = 2.0 is pure variance reduction; its disadvantage at λ_true = 3.0 is bias inflation.

Robustness of Behavioral Conclusions to Prior Choice

Table 2 presents P(λ > 2) from real IS posterior samples. At λ_true = 1.5, P(λ > 2) = 0.000 for all priors—complete agreement. At λ_true = 2.0, n = 100: P(λ > 2) = 0.890 (Prior A), 0.843 (Prior B), 0.762 (Prior C), range = 12.8 pp. At λ_true = 2.5, n = 100: P(λ > 2) = 0.989, 0.929, 0.835, range = 15.4 pp. Applying a 90% decision threshold, Prior B yields a positive conclusion but Prior C does not at λ_true = 2.5, n = 100—a qualitative discrepancy driven entirely by prior specification. This cross-prior disagreement showed minimal attenuation with sample size: at λ_true = 2.5, the range was 15.4 pp (n = 100), 8.1 pp (n = 200), and 8.1 pp (n = 500), confirming that increasing sample size did not resolve the directional conclusion divergence when λ_true was substantially above the canonical threshold.

Table 2. Posterior Directional Probability P(λ > 2) from 1,080 Bayesian IS Posterior

True λ	n	P(λ>2) A	P(λ>2) B	P(λ>2) C	Max Diff (pp)
1.5	100	0.000	0.000	0.000	0.0
1.5	500	0.000	0.000	0.000	0.0
2.0	100	0.890	0.843	0.762	12.8
2.0	200	0.764	0.737	0.703	6.1
2.0	500	0.602	0.598	0.573	2.9
2.5	100	0.989	0.929	0.835	15.4
2.5	200	0.995	0.969	0.914	8.1
2.5	500	0.996	0.967	0.915	8.1
3.0	100	0.990	0.929	0.840	15.0
3.0	200	0.996	0.968	0.917	7.9
3.0	500	0.996	0.967	0.914	8.2

Note. P(λ > 2) = posterior probability that λ exceeds the canonical threshold of 2.0, averaged over 30 replicates. Max Diff = maximum pairwise difference across priors (pp). All values from real Bayesian IS posteriors with ESS 577–832.

The directional probability formula underpinning Table 2 is:

$$P(\lambda \ > \ 2) = \ 1 - \Phi\left[\frac{(2 - \hat{\lambda}_{MAP})}{\sigma\_post}\right] \quad (8)$$

At λ_true = 2.5, n = 100, the real IS posteriors gave P(λ > 2) = 0.835 for Prior C versus 0.929 for Prior B—a 9.4 pp gap. A researcher using the 90% decision threshold would conclude loss aversion is above 2.0 under Prior B but not Prior C, despite both analyses using the same data. This qualitative discrepancy occurred in 100% of the 30 replicates at this condition, making it a systematic effect rather than a sampling artefact. The cross-prior disagreement at n = 500 (8.1 pp) being essentially the same as at n = 200 (8.1 pp) confirms that sample size is not an effective mitigation for prior-induced conclusion divergence in Bayesian prospect theory.

The cumulative evidence from Sections 3.1–3.5 confirms all three hypotheses stated in the Introduction. H1 was confirmed: Prior A (diffuse default prior) failed catastrophically for λ_true ≥ 2.0 due to prospect theory likelihood non-identifiability, with RMSE exceeding 3.5 across all sample sizes. H2 was confirmed: Prior C (Strong) introduced bias of −0.840 to −0.865 at λ_true = 3.0 that did not diminish from n = 100 to n = 500. H3 was partially confirmed: Prior B showed modest bias attenuation at λ_true = 2.0 (bias +0.329 to +0.131 from n = 100 to n = 500), but no attenuation for extreme λ_true values. Together, these results reframe prior sensitivity in behavioral science from a theoretical concern into a measured, model-specific, and practically consequential risk.

An unexpected finding, which was not anticipated in the a priori hypotheses, was that the diffuse default Prior A produced substantially larger estimation errors than the strongly informative Prior C when λ_true ≥ 2.0. For example, at λ_true = 2.0 and n = 100, the RMSE under Prior A (3.715) was 20.7 times larger than that observed under Prior C (0.179). This pattern suggests that parameter recovery in Bayesian Prospect Theory models may be particularly sensitive to prior specification when the likelihood surface provides limited information about the parameter of interest. Although previous studies have demonstrated the benefits of informative priors in Bayesian estimation (e.g., Smid et al., 2020), the present findings indicate that the magnitude of prior influence may be considerably larger in nonlinear behavioral decision models than has been documented in more conventional linear modeling contexts. This observation is broadly consistent with discussions of parameter identifiability in cognitive and behavioral models (Wilson & Collins, 2019).

From a broader methodological perspective, the present study extends previous work on prior sensitivity by examining a nonlinear Prospect Theory framework under controlled simulation conditions. Whereas earlier studies have often focused on linear or relatively well-identified models, the current findings highlight the possibility that prior influence may remain substantial even at larger sample sizes when model identifiability is limited. In this sense, the results should be viewed as complementary to existing recommendations regarding Bayesian practice rather than as contradictory evidence. More specifically, they suggest that conclusions about the diminishing influence of priors with increasing sample size may depend on the characteristics of the underlying model and likelihood function. The present findings therefore contribute to the growing literature on prior sensitivity by illustrating how model structure can shape the practical consequences of prior specification in behavioral decision research.

CONCLUSIONS AND SUGGESTIONS

This study demonstrated that prior specification is a statistically consequential methodological decision in Bayesian estimation of loss aversion under Prospect Theory, with the magnitude of prior influence depending systematically on the interaction between prior strength, true parameter value, and sample size. The simulation evidence established that weakly informative diffuse priors failed to identify the loss aversion parameter when the likelihood surface was flat, that strongly informative priors introduced persistent shrinkage bias that was not attenuated by increasing sample size, and that a theory-aligned moderate prior consistently achieved the most favourable balance between bias and estimation precision across all conditions tested. These findings confirm that the common assumption of prior irrelevance at moderate-to-large sample sizes does not generalise to nonlinear behavioral decision models, and that sensitivity analysis across prior specifications is a necessary component of rigorous Bayesian prospect theory research.

The principal limitation of this study is that it examined a single-level (non-hierarchical) model for a single behavioral parameter under a specific softmax choice rule. Future work should extend the framework to hierarchical prospect theory models that pool information across participants, to other behavioral parameters such as the probability weighting exponent and the temporal discount rate, and to real behavioral datasets from clinical populations where λ may deviate substantially from the population norm of 2.0. Integration with formal expert elicitation methods (Stefan et al., 2022; Gronau et al., 2020) would allow the construction of principled informative priors grounded in domain expertise, potentially further improving the balance between bias and precision demonstrated here for the moderate prior. These extensions would advance the goal of transparent, reproducible, and statistically rigorous Bayesian inference in behavioral science.

REFERENCE

Brown, A., Imai, T., Vieider, F., & Camerer, C. (2024). Meta-analysis of empirical estimates of loss aversion. Journal of Economic Literature, 62(2), 485–516. https://doi.org/10.1257/jel.20221698

Button, K. S., Ioannidis, J. P. A., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S. J., & Munafò, M. R. (2013). Power failure: Why small sample size undermines the reliability of neuroscience. Nature Reviews Neuroscience, 14(5), 365–376. https://doi.org/10.1038/nrn3475

Chakraborty, A., Nott, D. J., Drovandi, C. C., Frazier, D. T., & Sisson, S. A. (2022). Weakly informative priors and prior-data conflict checking for likelihood-free inference. arXiv:2202.09993. https://arxiv.org/abs/2202.09993

Depaoli, S., & van de Schoot, R. (2017). Improving transparency and replication in Bayesian statistics: The WAMBS-Checklist. Psychological Methods, 22(2), 240–261. https://doi.org/10.1037/met0000065

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). Bayesian data analysis (3rd ed.). CRC Press.

Gronau, Q. F., Ly, A., & Wagenmakers, E.-J. (2020). Informed Bayesian t-tests. The American Statistician, 74(2), 137–143. https://doi.org/10.1080/00031305.2018.1562983

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica, 47(2), 263–291. https://doi.org/10.2307/1914185

Nilsson, H., Rieskamp, J., & Wagenmakers, E.-J. (2011). Hierarchical Bayesian parameter estimation for cumulative prospect theory. Journal of Mathematical Psychology, 55(1), 84–93. https://doi.org/10.1016/j.jmp.2010.08.006

Nott, D. J., Wang, X., Evans, M., & Englert, B.-G. (2020). Checking for prior-data conflict using prior-to-posterior divergences. Statistical Science, 35(2), 234–253. https://doi.org/10.1214/19-STS731

Scheibehenne, B., & Pachur, T. (2015). Using Bayesian hierarchical parameter estimation to assess the generalizability of cognitive models of choice. Psychonomic Bulletin & Review, 22(2), 391–407. https://doi.org/10.3758/s13423-014-0684-4

Smid, S. C., McNeish, D., Miocevic, M., & van de Schoot, R. (2020). Bayesian versus frequentist estimation for structural equation models in small sample contexts: A systematic review. Structural Equation Modeling, 27(1), 131–161. https://doi.org/10.1080/10705511.2019.1577140

Stefan, A. M., Katsimpokis, D., Gronau, Q. F., & Wagenmakers, E.-J. (2022). Expert agreement in prior elicitation and its effects on Bayesian inference. Psychonomic Bulletin & Review, 29(5), 1776–1794. https://doi.org/10.3758/s13423-022-02074-4

Tierney, L., & Kadane, J. B. (1986). Accurate approximations for posterior moments and marginal densities. Journal of the American Statistical Association, 81(393), 82–86. https://doi.org/10.1080/01621459.1986.10478240

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. Journal of Risk and Uncertainty, 5(4), 297–323. https://doi.org/10.1007/BF00122574

van de Schoot, R., Depaoli, S., King, R., Kramer, B., Märtens, K., Tadesse, M. G., Vannucci, M., Gelman, A., Veen, D., Willemsen, J., & Yau, C. (2021). Bayesian statistics and modelling. Nature Reviews Methods Primers, 1(1), Article 1. https://doi.org/10.1038/s43586-020-00001-2

van Erp, S., Oberski, D. L., & Mulder, J. (2019). Shrinkage priors for Bayesian penalized regression. Journal of Mathematical Psychology, 89, 31–50. https://doi.org/10.1016/j.jmp.2018.12.004

Veenman, M., Stefan, A. M., & Haaf, J. M. (2023). Bayesian hierarchical modeling: An introduction and reassessment. Behavior Research Methods, 56(5), 4602–4620. https://doi.org/10.3758/s13428-023-02204-3

Walasek, L., & Stewart, N. (2015). How to make loss aversion disappear and reappear: A theoretical and empirical investigation of manipulations of loss aversion. Journal of Experimental Psychology: General, 144(1), 7–11. https://doi.org/10.1037/xge0000039

Wilson, R. C., & Collins, A. G. (2019). Ten simple rules for the computational modeling of behavioral data. eLife, 8, Article e49547. https://doi.org/10.7554/eLife.49547