Homework Assignment 3

Bayesian Learning

Professional Master in Economics, Insper

Instructor:	Prof. Hedibert F. Lopes
Teaching assistant/grader:	Guilherme Piantino
Issued date:	June 1st, 2026
Due date:	June 15th, 2026 (at most by 7:15pm)

Instructions

- Submit a single PDF (typeset in LaTeX; handwritten work will not be accepted) together with reproducible code in a single `.zip` archive.
- This assignment requires **R**, since it uses `rstan` (Hamiltonian Monte Carlo via Stan) and the `stochvol` package of Gregor Kastner and Darjus Hosszejni. You may use Python only for auxiliary data handling/plots, but the two estimation engines must be `rstan` and `stochvol`.
- All random-number experiments must be seeded so that results are reproducible. Use `set.seed(20260601)` (and pass an explicit `seed` to `rstan::sampling` and to the `stochvol` samplers) so that the class works from the same simulated streams.
- Effective sample sizes (ESS) and $\widehat{R}$ should be computed using standard tooling: `rstan::monitor` / `posterior` for the Stan output, and `coda` or the built-in diagnostics for the `stochvol` output. Model comparison should use the `loo` package (LOO/WAIC) where requested.
- Plots must be labeled and legible. Discussion of results is required throughout; raw numerical output without commentary will not receive full credit.
- The starting point for the Stan code is the worked example *"SV with Gaussian or Student's t errors"* available at `https://hedibert.org/wp-content/uploads/2025/05/sv-stan.html`. You are expected to *adapt and extend* that code; copying it verbatim without the requested extensions (leverage, heavy tails, model comparison) will not receive credit.
- Questions about the assignment should be directed to the TA, Guilherme Piantino.

The model: SV-AR(1)

Let $y_1, \dots, y_n$ be (mean-corrected) log-returns of a financial asset. The stochastic volatility model with an AR(1) log-variance is

$$y_t = \exp(h_t/2)\,\varepsilon_t, \qquad t = 1, \dots, n, \tag{1}$$

$$h_t = \mu + \phi\,(h_{t-1} - \mu) + \sigma\,\eta_t, \qquad t = 2, \dots, n, \tag{2}$$

$$h_1 \sim \mathcal{N}\left(\mu,\ \frac{\sigma^2}{1-\phi^2}\right), \tag{3}$$

with $|\phi| < 1$ (stationarity). The parameters are the level μ, the persistence ϕ, and the volatility-of-log-volatility $\sigma > 0$. The four model variants studied in this assignment differ in the joint distribution of the shocks (ε_t, η_t):

(i) **Gaussian, no leverage:** $\varepsilon_t \sim \mathcal{N}(0,1)$ independent of $\eta_t \sim \mathcal{N}(0,1)$.

(ii) **Student's t, no leverage:** $\varepsilon_t \sim t_\nu$ standardized to unit scale (heavy tails), independent of $\eta_t \sim \mathcal{N}(0,1)$; estimate the degrees of freedom ν.

(iii) **Gaussian, with leverage:** (ε_t, η_t) jointly normal with $\mathrm{corr}(\varepsilon_t, \eta_t) = \rho$, so that a return shock and the *contemporaneous* log-variance innovation are correlated; $\rho < 0$ is the usual "leverage" sign.

(iv) **Student's t with leverage:** heavy tails *and* leverage simultaneously (the full model).

In the `stochvol` package these correspond exactly to the wrappers `svsample` (vanilla), `svtsample` (t-errors), `svlsample` (leverage), and `svtlsample` (t-errors and leverage).

For the Stan implementation, use the *non-centered* parameterization of the states ($h_t = \mu + \sigma\, \tilde{h}_t/\sqrt{\cdot}$ with standardized AR(1) $\tilde{h}_t$) for better HMC geometry, and place priors comparable to the `stochvol` defaults, e.g.

$$\mu \sim \mathcal{N}(0, 10), \qquad \frac{\phi + 1}{2} \sim \mathrm{Beta}(20, 1.5), \qquad \sigma \sim \text{Half-Normal}(0, 1),$$

$$\nu \sim \mathrm{Gamma}(2, 0.1) \text{ truncated to } \nu > 2, \qquad \rho \sim \mathrm{Uniform}(-1, 1).$$

State clearly whichever priors you use and keep them consistent between the two engines so the comparison in part (f) is meaningful.

Question 1 (100 points)

Stochastic volatility for asset returns: Gaussian vs. Student's t errors and the leverage effect, estimated with Stan (HMC) and with `stochvol` (ASIS/interweaving).

This is a single, cumulative case study. You will build the base SV-AR(1) model and then add heavy tails and leverage, implementing every variant *twice* — once in `rstan` and once in `stochvol` — and finish with a head-to-head comparison.

(a) Data choice and exploratory analysis (10 points) Choose *one* return series of your own — a single stock, an equity index, an FX rate, or a cryptocurrency — with at least $n \geq 1{,}000$ daily observations. Download it reproducibly (e.g. via `quantmod`, `tidyquant`, or a CSV you commit to your `.zip`), state the ticker, source, and sample period, compute log-returns $y_t = \log(P_t/P_{t-1})$, and subtract the sample mean. Report:

- the time-series plot of y_t and basic summary statistics (mean, sd, skewness, excess kurtosis);
- the autocorrelation functions of y_t and of y_t^2 (or $|y_t|$), commenting on *volatility clustering*;
- a normal QQ-plot of y_t, commenting on heavy tails (motivation for the Student's t variant);
- a simple leverage diagnostic, e.g. the sample correlation between y_t and y_{t+1}^2 (or between the sign of y_t and future squared returns), commenting on whether a negative ρ is plausible.

(b) Base model in Stan: Gaussian, no leverage (20 points) Adapting the code at the link above, implement variant (i) in `rstan` using the non-centered parameterization. Run at least 4 chains with enough iterations to obtain bulk- and tail-ESS in the hundreds. Report:

- your Stan program (in an appendix) and the priors used;
- convergence diagnostics: $\widehat{R}$, number of divergent transitions, bulk/tail-ESS, and trace plots for (μ, ϕ, σ);
- marginal posteriors (mean, sd, 95% credible interval) for μ, ϕ, and σ;
- the posterior of the time-varying volatility $\exp(h_t/2)$ as a fan chart (posterior median and a 95% band) overlaid on $|y_t|$.

(c) Heavy tails in Stan: Student's t errors (15 points) Extend (b) to variant (ii) by giving ε_t a standardized Student's t_ν distribution and placing a prior on ν. Report the posterior of ν and discuss whether the data support heavy tails. Compare the estimated volatility path and the posteriors of (μ, ϕ, σ) to those of the Gaussian model. Compute the LOO (or WAIC) of models (i) and (ii) with the `loo` package and report the difference with its standard error.

(d) Leverage in Stan, and the full model (20 points) Extend (b) to variant (iii) by letting (ε_t, η_t) be bivariate normal with correlation ρ. Write down explicitly the conditional distribution of h_t given (h_{t-1}, y_{t-1}) that your implementation uses, and state your prior on ρ. Estimate ρ and interpret its sign and magnitude. Then implement variant (iv) (Student's t *and* leverage). Rank all four models by LOO/WAIC and state which is preferred for *your* series.

(e) The `stochvol` package (20 points) Re-estimate all four variants with Kastner and Hosszejni's `stochvol`, using `svsample`, `svtsample`, `svlsample`, and `svtlsample`, respectively (equivalently, the unified `svsample(...)` interface with the corresponding `priorspec`). Match the priors to those used in Stan as closely as the package allows (document any unavoidable differences). Report, for each variant, the posterior summaries of the parameters — including ν and ρ where present — and the estimated volatility path (`volplot`/`predict` or the posterior of $\exp(h_t/2)$).

(f) Comparison and synthesis (15 points) Produce a single master table comparing the **four models** $\times$ **two engines**. For each cell report the posterior mean and 95% credible interval of (μ, ϕ, σ) and, where applicable, ν and ρ; the sampling efficiency (minimum ESS per second of wall-clock time); the total runtime; and the LOO/WAIC for the Stan fits. Then discuss:

- Do the Stan and `stochvol` posteriors agree (overlay or tabulate the marginals)? Explain any discrepancies in terms of priors or parameterization.
- Does your series exhibit heavy tails (posterior of ν) and/or leverage (posterior of ρ)? Tie this back to the exploratory diagnostics in (a).
- Which engine is more efficient *per effective draw*, and why? Relate your answer to how `stochvol` samples the SV model (ancillarity–sufficiency interweaving, ASIS) versus how Stan explores the joint posterior of the $n + \mathcal{O}(1)$ parameters with HMC.
- Which of the four models would you report for your asset, and why? One or two paragraphs.

References

- Kastner, G. and Frühwirth-Schnatter, S. (2014). Ancillarity-sufficiency interweaving strategy (ASIS) for boosting MCMC estimation of stochastic volatility models. *Computational Statistics & Data Analysis* 76, 408–423.
- Hosszejni, D. and Kastner, G. (2021). Modeling univariate and multivariate stochastic volatility in R with `stochvol` and `factorstochvol`. *Journal of Statistical Software* 100(12).
- Gamerman, D. and Lopes, H. F. (2006). *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*, 2nd ed., Chapman & Hall/CRC.
- Lopes, H. F. (2025). *SV with Gaussian or Student's t errors* (Stan/RStan worked example). `https://hedibert.org/wp-content/uploads/2025/05/sv-stan.html`
- Stan Development Team. *Stan User's Guide* — Stochastic Volatility chapter.

End of assignment.