A Appendix

A.1 Proof of Proposition 8

We establish a monotonicity property of the proposed coupling for PIMH. Monotonicity properties of IMH samplers were exploited in an exact simulation context in [8] without this explicit construction.

Proposition 10. *Under the proposed coupling scheme the sequence of likelihood estimates* $(p_N^{(n+1)})_{n\geq 0}$ *stochastically dominates* $(\tilde{p}_N^{(n)})_{n\geq 0}$ *in the sense that for any* $n \geq 1$, $s \geq 0$

$$p_N^{(n)} \geq \tilde{p}_N^{(n-1)} \Rightarrow p_N^{(n+s)} \geq \tilde{p}_N^{(n+s-1)} \quad a.s.$$

Proof. The coupling procedure in Algorithm 3 uses a single proposal p_N^* and samples $\mathfrak{u} \sim \mathcal{U}[0,1]$, with proposals being accepted according to:

$$\text{if } \mathfrak{u} \leq 1 \wedge \frac{p_N^*}{p_N^{(n)}} \text{ then } p_N^{(n+1)} = p_N^*, \text{ else } p_N^{(n+1)} = p_N^{(n)},$$

$$\text{if } \mathfrak{u} \leq 1 \wedge \frac{p_N^*}{\tilde{p}_N^{(n-1)}} \text{ then } \tilde{p}_N^{(n)} = p_N^*, \text{ else } \tilde{p}_N^{(n)} = \tilde{p}_N^{(n-1)}.$$

We see that if $p_N^{(n)} \geq \tilde{p}_N^{(n-1)}$ then either

1. $p_N^{(n+1)} = p_N^*$ in which case
$$\mathfrak{u} \leq 1 \wedge \frac{p_N^*}{p_N^{(n)}} \implies \mathfrak{u} \leq 1 \wedge \frac{p_N^*}{\tilde{p}_N^{(n-1)}}$$
as $p_N^{(n)} \geq \tilde{p}_N^{(n-1)}$ so that $p_N^{(n+1)} = \tilde{p}_N^{(n)} = p_N^*$ (i.e. the chains meet).
2. $p_N^{(n+1)} = p_N^{(n)}$ and so $p_N^* \leq p_N^{(n)}$. In this case, either $\tilde{p}_N^{(n)} = p_N^*$, and so $p_N^{(n+1)} \geq \tilde{p}_N^{(n)}$, or $\tilde{p}_N^{(n)} = \tilde{p}_N^{(n-1)}$ in which case both chains have rejected p_N^* and the ordering is preserved.

Finally, from the initialization of the procedure, we have $p_N^{(1)} \geq \tilde{p}_N^{(0)}$ because the initial state of the second chain is used as a proposal in the first iteration of the first chain. □

From the above reasoning we see that the chains meet when the first chain accepts its proposal for the first time, as the second chain then necessarily accepts the same proposal.

From the initial state with likelihood estimate $p_N^{(0)}$, the acceptance probability of the first chain is $\int 1 \wedge (p_N/p_N^{(0)})\bar{g}(p_N)\mathrm{d}p_N$, with $\bar{g}$ denoting the density of the PF likelihood estimator p_N. Thus, the time to the first acceptance follows a Geometric distribution with success probability $\int 1 \wedge (p_N/p_N^{(0)})\bar{g}(p_N)\mathrm{d}p_N$. The result stated in Proposition 8 follows when rewriting the problem using the error of the log-likelihood estimator $\log\{p_N(y_{1:T})/p(y_{1:T})\}$.

A.2 Integrated autocorrelation time for various test functions

We show here experimentally that IF(h) is approximately proportional to IF(σ) for various test functions: $h_1 : x_{1:T} \mapsto x_1$, $h_2 : x_{1:T} \mapsto x_T$, $h_3 : x_{1:T} \mapsto \sum_t x_t$ and $h_4 : x_{1:T} \mapsto \sum_t x_t^2$. This is illustrated in Figure 7 where IF(h) is displayed for a range of N against IF(σ) over the corresponding range of σ.

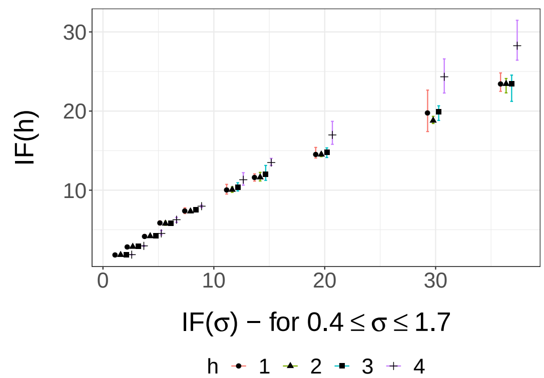

Figure 7: Inefficiency IF$[h_i]$ versus IF$[\sigma]$, with markers indicating the test functions h_1, h_2, h_3, h_4. The vertical axis scale is relative, depending on the test function.

A.3 Rao-Blackwellisation for stochastic kinetic model

We demonstrate here the gains arising from the use of a Rao-Blackwellized estimator detailed in Section 2.2. We display in Figure 8 the variance of the two unbiased estimators of $\mathbb{E}(X_{1,\Delta t}|y_{1:T})$ for $t = \Delta, ..., T\Delta$ and $T = 100$ for the latent Markov jump process and two different values of N, 200 and $3,000$. In both cases we set $k = m = 0$. As expected $\bar{H}_{0:0}$ outperforms

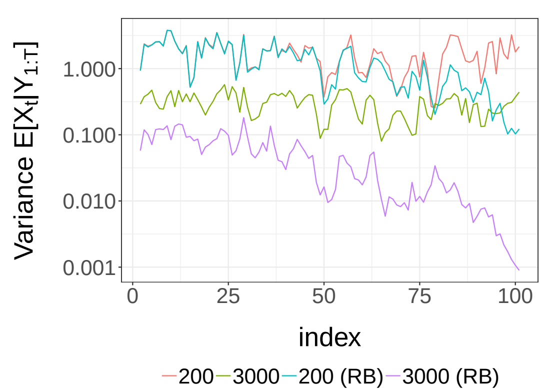

Figure 8: Empirical variance of unbiased estimators of $\mathbb{E}(X_{1,\Delta t}|y_{1:T})$: $H_{0:0}$ and Rao-Blackwellised (RB) estimator $\bar{H}_{0:0}$ for stochastic kinetic model.

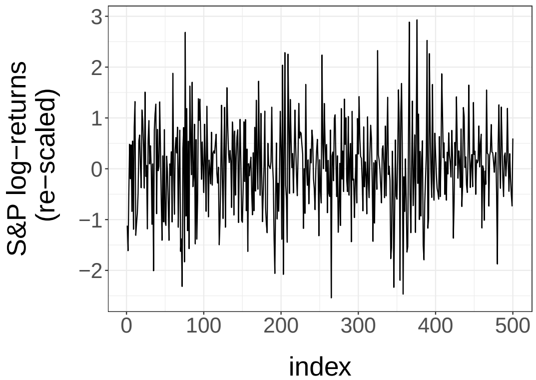

(a) Daily returns of S&P 500 data

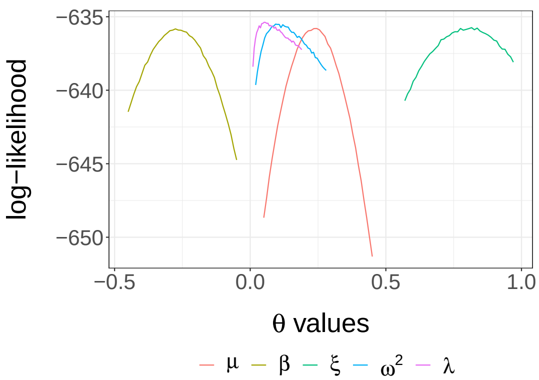

(b) Particle estimates of the log-likelihood function

Figure 9: Data and parameter estimation for Lévy-driven stochastic volatility model.

$H_{0:0}$ but the benefits are much higher for t close to T than when t is close to 1. For example, we see that for $N = 200$ the estimators $H_{0:0}$ and $\bar{H}_{0:0}$ coincide for $t \leq 35$. This is an expected consequence of the particle path degeneracy problem [13, 23, 25], with many particles $(X^i_{1:T})_{i\in[N]}$ obtained by the PF at time T sharing common ancestors for t close to 1 when N is too small; see [23] for results on the corresponding coalescent time.

A.4 Data and parameter estimation

The data used comprised of $T = 500$ log-returns of the S&P 500, starting from the 3^{rd} January 2005, with the scaled data taken from the stochastic volatility example used to demonstrate SMC^2 in [6]. We plot the raw data in Figure 9a. Parameters were estimated using a two-stage procedure, with SMC^2 used to find a region of high marginal likelihood under the model. Further refinement was performed to compute the maximum likelihood estimator (MLE) $\hat{\theta}$ of $\theta = (\mu, \beta, \xi, \omega^2, \lambda)$ using a grid search around values close to the optimum using $N = 10,000$ particles. We obtained $\hat{\theta} = (0.24, -0.28, 0.82, 0.09, 0.05)$. Likelihood curves around the optimal values are shown in Figure 9b, where for each parameter component the log-likelihood was varied while keeping the other parameters fixed at $\hat{\theta}$.

B Application of unbiased estimation to SMC samplers

SMC samplers are a class of SMC algorithms that can be used in Bayesian inference to approximate expectations w.r.t. complex posteriors for static models [10]. We show here how we can directly use the methodology proposed in this paper to obtain unbiased estimators of these expectations.

B.1 Bayesian computation using SMC samplers

Assume one is interested in sampling from the posterior density $\pi(x) \propto \nu(x)L(x)$ where $\nu(x)$ the prior density w.r.t. a suitable dominating measure and $L(x)$ is the likelihood. We also assume that one can sample from ν. To approximate π, a specific version of SMC samplers introduces a sequence of $T-1$ intermediate densities π_t for $t = 2, ..., T$ bridging ν to π using

$$\gamma_t(x) = \nu(x)L^{\beta_t}(x), \qquad \pi_t(x) = \frac{\gamma_t(x)}{\mathcal{Z}_t},$$

where $\beta_1 = 0 < \beta_2 < ... < \beta_T = 1$. The choice of the sequence $\{\beta_t : t = 2, ..., T-1\}$ can be guided using a preliminary adaptive SMC scheme, as in [33], which should subsequently be fixed to preserve unbiasedness of the normalizing constant estimate and validity of the resulting PIMH. In SMC samplers, particles are initialized at time $t = 1$ by sampling from the prior ensuring $w_1(x_1) = 1$. At time $t \geq 2$, particles are sampled according to an MCMC kernel leaving π_{t-1} invariant and are then weighted according to

$$w_t(x_{t-1}, x_t) = \frac{\gamma_t(x_{t-1})}{\gamma_{t-1}(x_{t-1})} = L^{\beta_t - \beta_{t-1}}(x_{t-1}).$$

Particles are resampled according to these weights and we set

$$\mathcal{Z}_{t,N} = \mathcal{Z}_{t-1,N} \cdot \frac{1}{N}\sum_{i=1}^{N} w_t(X_{t-1}^{A^i_{t-1}}, X^i_t),$$

with $\mathcal{Z}_{1,N} = 1$. At time T, $\pi_N(\mathrm{d}x_T) := \sum_{i=1}^N W^i_T \delta_{X^i_T}(\mathrm{d}x_T)$ provides a Monte Carlo approximation of the distribution π and $\mathcal{Z}_{T,N}$ plays the role of $p_N(y_{1:T})$, approximating the normalizing constant $\mathcal{Z}_T$ of $\pi = \pi_T$. If no resampling is used, this specific version of SMC samplers coincides with AIS [30] in which case $\mathcal{Z}_{T,N}$ is given by the average of the product of the incremental weights from time $t = 1$ to

$t = T$ instead of the product of the averaged incremental weights. We can use this SMC sampler algorithm or AIS directly within the coupled PIMH scheme, replacing $p_N(y_{1:T})$ by $\mathcal{Z}_{T,N}$ in the acceptance probabilities. We see that $\sup_{(x,x')\in \mathsf{X}^2} w_t(x, x') < \infty$ provided that $\sup_{x\in\mathsf{X}} L(x) < \infty$. Under this condition, if Assumption 5 is satisfied then the estimator $\bar{H}_{k:m}$ of $\pi(h)$ is unbiased and has finite variance and finite expected cost.

B.2 Numerical example

We use here coupled PIMH to debias expectations w.r.t. the posterior distribution for a Bayesian mixture model discussed in [28]. We have

$$L(x) = \prod_{n=1}^{M} \left(\frac{1}{D} \sum_{i=1}^{D} \mathcal{N}(y_n; x^i, \sigma^2) \right)$$

with $x := (x^1, ..., x^D) \in \mathbb{R}^D$ constituting the unknown mean components. We consider here $D = 4$ mixture components and $M = 100$ observations. A uniform prior distribution is placed on x over the hypercube $[-10, 10]^D$. The resulting posterior distribution is multimodal. We set $\sigma = 1$ and simulate observations from the model with true values $x^* = (-3, 0, 3, 6)$. We adopt a symmetric random walk for the Metropolis–Hastings proposals with identity covariance and pick $\beta_t = \left(\frac{t-1}{T-1}\right)^2$ for $T = 200$.

We simulate 10,000 estimators with $m = 64$ and $N = 100$, after which we are able to estimate variance of test functions for a range of values of k and m noting that there will be some correlation introduced between estimators. The results are shown in Figure 10 where we plot the meeting times of the unbiased estimators in Figure 10a and the variance of the estimators for a range of values of m in Figure 10b using $h : x \mapsto x^1 + x^2 + (x^1)^2 + (x^2)^2$. Uncertainty in the estimated values of $\mathbb{V}[\bar{H}_{k:m}(h)]$ was obtained using 1,000 bootstrap samples, resampling 10,000 of the 10,000 unbiased estimators with replacement. The figure shows the variance of these estimators for a range of values of $m \in \{4, 8, ..., 64\}$ while varying $k \in \{0, ..., m-1\}$. We see, firstly, as expected that as m increases the variance of the estimators reduces. Secondly, for each value of m we see that there exists an optimal value of k, however, as m increases the optimum becomes less pronounced, suggesting that as m increases there is a degree of insensitivity to the choice of k. Finally, for the range of m considered, the optimal values of k appear in a comparatively small interval close to the origin, suggesting that it is not necessary to use large values of k to reduce the variance contribution arising from the bias correction.

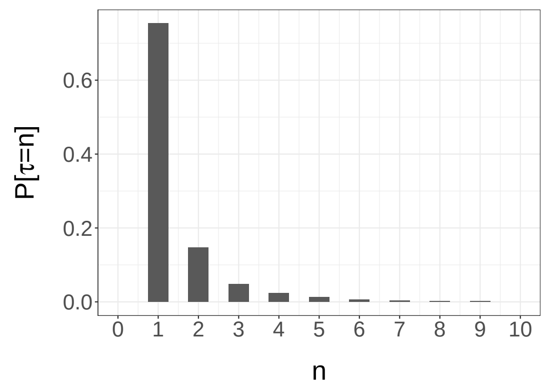

(a) Empirical distribution of meeting times for SMC sampler with $N = 100$ over 10,000 independent runs. The estimated 95^{th} and 99^{th} percentiles were 6 and 13 respectively.

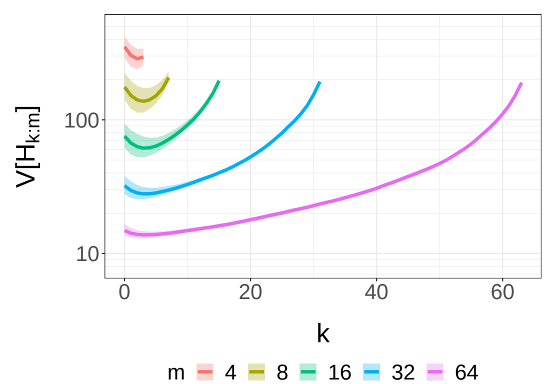

(b) $\mathbb{V}[\bar{H}_{k:m}]$ of Rao-Blackwellised unbiased estimators of $\pi(h)$ for SMC sampler as a function of k for a range of values of m. The shaded regions correspond to the 1^{st} and 99^{th} percentiles of the variance estimator.

Figure 10: SMC sampler unbiased estimators