The left tail methods of Asmussen, Ben Amar and Ben Rached

Taking stock of Asmussen (2016), Ben Amar (2023) and Ben Rached (2024)

Stephen J. Mildenhall

2026-10-22 19:55 commit b5ca906+

1 Summary of Asmussen et al. (2016), Ben Amar et al. (2023), Ben Rached et al. (2024)

The specific test problems are the weak part. "Outage probability of a diversity receiver in log-normal fading" is a canonical benchmark that gets reused paper after paper, and for the parameter ranges actually tabulated, a careful convolution on a laptop gets four significant figures in seconds. Asmussen's entire apparatus, Lambert W asymptotics for the Laplace transform, an approximate Cramér function, acceptance-rejection from the tilted density, 100,000 replications, reproduces in five build calls. That is a comment on the choice of benchmark, not on the mathematics.

There is also a spurious-precision problem underneath. Ben Amar computes a 2×10^{-12} outage probability to 5% relative accuracy. Nobody knows a fading distribution to twelve significant figures. The model error exceeds the numerical error by fifty orders of magnitude, so the accuracy being competed over is not accuracy about anything.

But the methods are not nonsense, and the 2024 paper is the proof. Ben Rached, Hoel and Meo are the ones who broke ranks and asked "why not just convolve it", and in doing so produced the sharpest result of the three: a theorem about when convolution fails, which turns out to be a precise and checkable statement about `aggregate`'s own numerical floor. It located that floor for the first time. That is a real contribution, and it came out of the cottage industry rather than despite it.

The honest line is that importance sampling and saddlepoint methods earn their keep where grids genuinely die: dependence between summands, dimension above one, summands on wildly different scales, and tails below double precision. The published test problems exercise none of those, which is why they look easy. If you want the papers to look hard, correlate the log-normals.

The framing that fits is that these are three papers about the same problem in which `aggregate` is the referee rather than the competitor. It confirms Asmussen's second-order saddlepoint, catches a bad row in his table, lands on Ben Amar's seven quoted values, and then fails exactly where Ben Rached's theorem says it must.

Asmussen, S., Jensen, J. L., & Rojas-Nandayapa, L. (2016). Exponential Family Techniques for the Lognormal Left Tail. *Scandinavian Journal of Statistics*, *43*(3), 774–787. https://doi.org/10.1111/sjos.12203

Ben Amar, E., Ben Rached, N., Haji-Ali, A.-L., & Tempone, R. (2023). State-dependent importance sampling for estimating expectations of functionals of sums of independent random variables. *Statistics and Computing*. https://doi.org/10.1007/s11222-022-10202-2

Ben Rached, N., Hoel, H., & Meo, J. V. (2024). A fast and accurate numerical method for the left tail of sums of independent random variables. *Statistics and Computing*. https://doi.org/10.1007/s11222-024-10514-5