

A tale of two time scales: Determining integrated volatility with noisy high frequency data.

Per Mykland, University of Chicago

Computational Finance Seminar, March 24

Abstract

In the analysis of high frequency financial data, a major problem concerns the nonparametric determination of the volatility of an asset return process. A common practice is to estimate volatility from the sum of the frequently-sampled squared returns. Though this approach is justified under the assumption of a continuous stochastic model, it runs into the challenge from market microstructure in real applications. We argue that this customary way of estimating volatility is flawed in that it overlooks what is, in fact, observation error. The usual mechanism for dealing with the problem is to throw away some data, by sampling less frequently. We propose here a statistically sounder device. Our device is model-free, it takes advantage of the rich sources in tick-by-tick data, and to a great extent it corrects the effect of the microstructure noise on volatility estimation. We also develop a limit theory for the corrected estimator.