Economics Modeling with HPC

Nathaniel Throckmorton (W&M Economics)

My research studies the policies of central banks in a low interest rate environment and their effect on macroeconomic variables such as gross domestic product (GDP) and the inflation rate. At the core of that research are mathematical models of the macroeconomy that make equilibrium predictions about movements in those variables over time. The models, known as dynamic stochastic general equilibrium (DSGE) models, were born out of the 1970s rational expectations revolution in macroeconomics and have grown into a workhorse for macroeconomists at major universities and central banks, such as the Federal Reserve (the Fed).

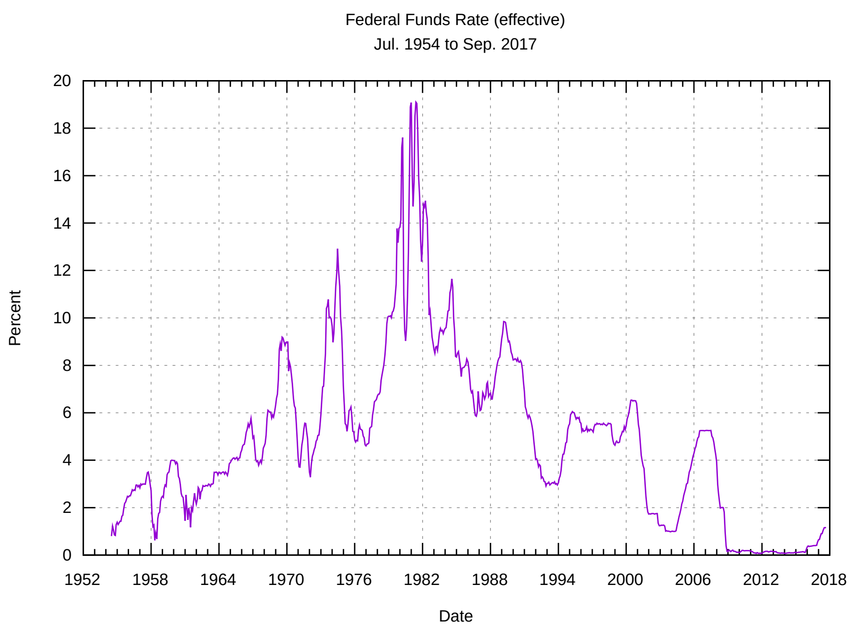

Figure 1: Federal Funds rate vs. time (from Wikipedia: "Interest_rate"). Notice the rate is quite close to zero from 2009 through 2015.

The consensus view is that central banks should respond to recessionary pressures (i.e., falling GDP and inflation rates), by cutting their target interest rate (policy rate). In practice, DSGE models are relatively easy to work with if the central bank can lower its policy rate during a recession. However, if a recession, such as the last one, is severe enough that they cannot lower their policy rate further, then DSGE models can make qualitatively and quantitatively different predictions about GDP and inflation. For example, the Fed responded to the 2008 global financial crisis by lowering its policy rate to the zero lower bound (ZLB), where it remained through the end of 2015 (see Figure 1). My research examines the consequences of the ZLB on a central bank's policy rate. To continue using DSGE models for policy analysis, macroeconomists need to utilize different and more expensive computational tools that account for the ZLB and its effects on a model's predictions.

Atkinson et al. (2019, forthcoming in the Journal of Monetary Economics) compares the performance of two popular estimation methods for DSGE models with a ZLB using artificial datasets. The goal of the paper is to inform researchers as to the benefits and costs associated with each method so they can decide which is best suited to their application. Researchers can either estimate a fully nonlinear model that accounts for precautionary savings effects of the ZLB or a piecewise linear model that is much faster but ignores the precautionary savings effects. We find the predictions of the nonlinear model are typically more accurate than the piecewise linear model, but the differences are usually small. There are far larger gains in accuracy from estimating a richer, less misspecified piecewise linear model.

In order to quantify the differences between the linear and nonlinear model, a large number of calculations were performed. Since the nonlinear models were quite expensive to compute, this project took a large number of cycles on our local W&M HPC cluster (using ~6.5 million processor-hours during this project) as well as resources at a few other universities / supercomputing centers across the country. The code for this project was originally written in MATLAB but was converted over to a parallel MPI code for the production calculations.

This work was reported by the *Flat Hat* in April of 2019:
https://www.wm.edu/news/stories/2019/when-your-research-gets-really-computational,-head-for-wms-giant-abacus.php

Citation for this work:
@Article{AtkinsonEtAl2019,
author = {Tyler Atkinson and Alexander W. Richter and Nathaniel A. Throckmorton},
title = {The Zero Lower Bound and Estimation Accuracy},
journal = {Journal of Monetary Economics},
year = {2019},
volume = {forthcoming},
doi = {https://doi.org/10.1016/j.jmoneco.2019.06.007}