FEDERAL RESERVE BANK of NEW YORK Serving the Second District and the Nation

Liberty Street Economics

« Just Released: Job Market Remains Tight as Regional Economy Slows | Main | At the New York Fed: Conference on the Evolving Structure of the U.S. Treasury Market »

DECEMBER 03, 2015

The FRBNY DSGE Model Meets Julia

Marco Del Negro, Marc Giannoni, Pearl Li, Erica Moszkowski, and Micah Smith

We have implemented the FRBNY DSGE model in a free and open-source language called Julia. The code is posted here on GitHub, a public repository hosting service. This effort is the result of a collaboration between New York Fed staff and folks from the QuantEcon project, whose aim is to coordinate development of high performance open-source code for quantitative economic modeling.

You may wonder why we wrote our code, which was originally in MATLAB and made available here, in Julia. MATLAB is a widely used, mature programming language that has served our purposes very well for many years. However, Julia has two main advantages from our perspective. First, as free software, Julia is more accessible to users from academic institutions or organizations without the resources for purchasing a license. Now anyone, from Kathmandu to Timbuktu, can run our code at no cost. Second, as the models that we use for forecasting and policy analysis grow more complicated, we need a language that can perform computations at a high speed. Julia boasts performance as fast as that of languages like C or Fortran, and is still simple to learn. (Read this post, written by the creators of the language, to understand why Julia fits the bill.) We want to address hard questions with our models—from understanding financial markets developments to modeling households' heterogeneity—and we can do so only if we are close to the frontier of programming.

We tested our code and found that the model estimation is about ten times faster with Julia than before, a very large improvement. Our ports (computer lingo for "translations") of certain algorithms, such as Chris Sims's gensys (which computes the model solution), also ran about six times faster in Julia than the MATLAB versions we had previously used. (These results should not be interpreted as a broad assessment of the speed of Julia relative to MATLAB, as they apply only to the code we have written.) This document written by the New York Fed and QuantEcon collaborators, who did the real work on the port, documents the speed improvements and offers an overview of the hurdles encountered in the translation of a large codebase from one language to another. We hope it will be of use to other central banks and researchers embarking on similar projects.

We posted our code on GitHub because it is a natural home for open-source projects like ours. Anyone can easily download the code and—most importantly from our point of view—offer suggestions on how to improve it by posting enhanced versions and extensions of our routines. This release also provides an opportunity for the research community to experiment with an open-source, large-scale dynamic stochastic general equilibrium (DSGE) model that is actively used in a research and policy setting. The

point of collaborative programming goes beyond the joy of sharing; it is a form of "model validation." We constantly test the accuracy of our code (and the process of translation into Julia led to yet another line-by-line examination), but we also believe that the best way of making sure that the code is accurate is by letting the rest of the world be our reviewers. If there are inaccuracies or inefficiencies, somebody will find them. If there is a way to make the code faster, somebody will suggest it.

Finally, we want to thank our friends from QuantEcon, Zac Cranko, Spencer Lyon, and Pablo Winant, who worked elbow to elbow on the code with our staff, as well as John Stachurski, who made the collaboration possible. We are not done with this project: what we posted was just the DSGE model estimation part of the code, and a lot more is yet to come. Check out our code, and stay tuned!

Disclaimer

The views expressed in this post are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

This post reflects the experience of the authors with Julia and MATLAB and does not represent an endorsement by the Federal Reserve Bank of New York or the Federal Reserve System of any particular product or service.

Marco Del Negro is an assistant vice president in the Federal Reserve Bank of New York's Research and Statistics Group.

[]

Marc Giannoni is an assistant vice president in the Bank's Research and Statistics Group.

[] Pearl Li is a undergraduate at the Wharton School of the University of Pennsylvania, and was a summer intern in the Group during the summer of 2015.

Erica Moszkowski is a research analyst in the Group.

Micah Smith is a senior research analyst in the Group.

Posted by Blog Author at 07:00:00 AM in DSGE, Macroecon

The comments to this entry are closed.