

Ramblings on Tractability in Macroeconomics

Fabio Ghironi[†]

University of Washington, CEPR, and NBER

Director, Research Program in International Trade and Macroeconomics, Central Bank Research Association

August 6, 2018; Updated on August 13, 2018

There was an interesting debate on Twitter and in the blogosphere last April on how the quest for tractability has been shaping the development of macroeconomic models. Most of the emphasis was on the costs of tractability.¹ I tweeted back then that I would write my thoughts and post them, but then the reality of the teaching quarter and too many tasks for my own good became a binding constraint, so I get to this only now. I think of myself as an applied macro theorist, in the sense that my primary interest is in the development of models to study (what I consider) empirically relevant mechanisms that can matter for policy analysis. Thus, I offer a view of tractability from the model-building trenches in this note.

From my perspective, tractability is the requirement that I must be able to understand transparently the mechanisms and results of the model at hand—whether it is a model that I am writing (or reading) for research purposes or a model that I teach. If I (and readers, and students) cannot understand where its results are coming from, the fact that computational improvements made it possible to solve the model does nothing for me as far as tractability is concerned.

For this reason, when I advise grad students who want to work on models, I always ask them to start from the simplest, most transparent version of the model they have in mind. I always emphasize the desirability of going as deep as possible into the machine with pencil and paper before turning to the computer—even when it is impossible in the end to solve the model fully with pencil and paper. I want my students to get dirty with the grease of the machine all the way to their shoulders (and beyond) by using tools (other than a keyboard) with their hands before starting to sit in front of a computer screen staring at numbers and impulse responses.

This has spawned interesting myths about me and computing tools—for instance, that I hate Dynare. I don't. But it is a double-edged sword. Here is an example of what I mean: Years ago, one of my students came to my office with something like a 4x6 printout of impulse responses (4 rows, 6 columns of diagrams) and told me: “I really have no idea what is going on here, but my goal was to put the model into Dynare by this date.” I counted to 10 (probably more), took a few deep breaths, and proceeded to tell the student that “No, your goal shouldn't be that. It should be to have a model you understand.” There are obviously other goals that we discussed, but I hope the point is clear: By making it easy to solve complicated models computationally, Dynare and other similar tools have resulted in the production of a plethora of papers where authors clearly did not spend

[†] Department of Economics, University of Washington, Savery Hall, Box 353330, Seattle, WA 98195, U.S.A. E-mail: ghiro@uw.edu. URL: <http://faculty.washington.edu/ghiro>. The views in this note are personal and do not reflect the views or policies of the CEPR, NBER, or the Central Bank Research Association.

¹ See, for instance, [this post](#) by Beatrice Cherrier and her follow-up post [here](#). [This older post](#) by Nick Rowe is also very interesting.

enough time getting their hands dirty and understanding the nuts and bolts of their machines. That's not what tractability is about as far as I am concerned.

Tractability is in building models that make it possible to shine a bright light on the mechanisms the model-builder wants to focus on. At times, this can be done with models that are simple enough to be solved completely with pencil and paper. At times it cannot. And yet, if sufficient time is spent removing unnecessary ingredients and getting one's hands dirty with the grease of the machine, even a model for which the computer is ultimately necessary can be tractable, in the sense of delivering results that can be understood and explained.

Here is another example from personal experience: My most successful work so far has been my 2005 *QJE* paper with Marc Melitz, in which we developed a general equilibrium, dynamic version of the Melitz trade model in the presence of aggregate uncertainty. Marc and I started brainstorming on that paper in the early spring of 2001. We wrote the first draft between May and June of 2003. Besides converging on a common language², Marc and I spent a huge amount of time figuring out how we could perform the task of merging his trade model into standard international macro in what we thought would be the most parsimonious and transparent way possible. We also spent a large amount of time working with our model with pencil and paper in the effort of sharpening our understanding of its results, even if the full solution of the model required us to turn to the computer. It was that effort that delivered the equation for the real exchange rate—equation (5) in our paper—that made it possible for us to understand fully the implications of our model for the Harrod-Balassa-Samuelson evidence and to explain them clearly to our audiences and readers. Without that effort, forget about *QJE* publication and the success the paper had—even if computing tools made it feasible to spit out impulse responses and tables of numbers.

Tractability is also tied to the purpose for which I am writing a model, and the model itself is tied to the purpose for which it is written. If what I want to do is to show how one specific mechanism works and how it can, say, contribute to shaping fluctuations, I will try to strip out of the model everything that is not directly relevant to that mechanism, even if that means omitting a lot of stuff that is empirically relevant for other purposes. An example of this approach in my work is the paper that Florin Bilbiie, Marc Melitz, and I published in the *JPE* in 2012 (BGM below). We wanted to study how producer entry subject to sunk costs could propagate and shape business cycles, and how different forms of consumer preferences could matter. We kept our model as simple as possible to accomplish our goal.³

If I am writing a model to perform an exercise that is supposed to provide guidance for a policy discussion (for instance, the debate on structural reforms of product and labor markets), I will write a model that incorporates features of reality that are not in BGM: There will be unemployment, capital in the production function will not be an “extension,” there will be adjustment costs, and

² I discovered that trade economists and macroeconomists often have different things in mind when referring to real variables, and it makes for some “funny” moments: “What exactly did you mean with real variables for the last hour? Let's not do like [NASA with its centimeters to inches Mars-probe screw-up!](#)”

³ Tractability is also related to solution technique: There are plenty of circumstances in which log-linearization is a bad idea, but there are also many scenarios in which it remains a valuable tool. If I am exploring a situation in which solving a model through non-linear methods ultimately delivers results that are very similar to those produced by log-linearization, and if I can manipulate the log-linear system with pencil and paper to understand and show transparently what is going on, then, hell yes to log-linearization—and put the figures from solving the model non-linearly in appendix if anyone asks for them!

more. Example? The work that I did with Matteo Cacciatore, Romain Duval, and Giuseppe Fiori that provided the theoretical background for the analysis of structural reforms in the April 2016 IMF *World Economic Outlook*. (See our *JEDC* 2016 article and our working paper on reforms at the zero lower bound.)

Different types of tractability: The BGM model is tractable enough that versions of it can literally be solved with pencil and paper; the work with Matteo, Romain, and Giuseppe is more complicated, but it builds on the understanding of each individual ingredient that we developed over time by working on them separately, and thus it preserves our ability to explain results transparently. Somewhere in between is the 2016 *JIE* paper that I wrote with Matteo and Giuseppe on structural reforms and optimal monetary policy in a monetary union: The model has more moving parts than BGM, but it is still simple enough that clear insights can be obtained through pencil-and-paper work before moving on to the computer.

These thoughts should also clarify that models, model development, and tractability are crucially tied to the purpose for which a model is written. And the previous paragraphs should at least suggest (and I will now make that explicit) that asking a model to do more than what it was written for can be a dangerous exercise: It should be done with caution and awareness that we may be pushing the model far outside the boundaries of what it was intended for.

Is it a damning flaw if a model doesn't do well when we ask it to do something different from what it was designed for? I don't think so, at least not necessarily. I'll explain this with an example. I find Barry Eichengreen's (1984 *EEH*) model of policy interactions under the interwar Gold Standard a beautiful tool to explain to students why non-cooperative monetary policymaking could result in a contractionary bias in monetary policy that contributed to the dynamics of the Great Depression. It is a beautifully tractable, pencil-and-paper model. I would not dream of using it for all sorts of purposes other than what it was written for. Does that make it any less valuable? I really don't think so. It is a great analytical tool to help flash a spotlight on mechanisms that were at work in the interwar years, as Eichengreen's discussion of historical events clarifies. A beautifully tractable, simple model that helps me understand the reality of historical events. No blogger, physicist-with-econ-envy, heterodox economist, or member of your favorite other category of mainstream econ bashers will ever convince me that I should ditch that model and stop using it for my teaching because, say, it does not forecast accurately time series X, Y, or Z, or because it makes some draconian simplifying assumptions.

By the way, the last paragraph, should also make it clear that there should be no commitment or "marriage" to any specific set of tools in the pursuit of tractable model building. Yes, I have been using DSGE for most of my work since graduate school. I explained my thoughts on why I think many standard criticisms of DSGE macro are misplaced, and why DSGE tools remain (and will remain) very valuable [here](#), [here](#), and [here](#). That does not mean that I think DSGE should be the only admissible tool. Far from it. I greatly enjoyed working on non-microfounded models of policy interactions in the spirit of Canzoneri and Henderson (1991 MIT Press) with Barry Eichengreen and Francesco Giavazzi during my grad school years and shortly after (and I still think of the question I studied in my undergrad thesis with a three-country model of fiscal and monetary policy interactions as one of the most interesting questions I ever studied⁴). I use non-microfounded

⁴ The question was how EMU would affect incentives for policy coordination between the United States and Europe. I developed a framework in which the model-U.S. (with its central bank and its government) interacted with a two-

tools in my teaching (and I ask graduate students who want to be serious about international macro to audit my undergraduate class so they see that material in addition to the microfounded models I teach in my graduate classes—if they don't audit the undergrad class, they have to read the non-microfounded stuff anyway to prepare for my part of the field exam). There are a lot of beautifully tractable models written in the past, long before the advent of DSGE, that we can learn a lot from. And there is a lot of non-DSGE model-based work going on now that yields important results and insights. It is great when these different types of work are tractable in the sense of transparent in their mechanisms, so we can all understand where the results (and, possibly, their differences when different types of models are applied to the same question) are coming from.

Scholars I deeply admire, and from whose work and teaching I learned much of what I know, can “see through the cloud” of details of complicated models, and quickly identify key mechanisms and intuitions because they spent large portions of their careers working through the simpler models—the pencil-and-paper variety—that now are building blocks of the more complicated setups. Take Maury Obstfeld, for instance, or Giancarlo Corsetti and Paolo Pesenti, or Mick Devereux: These are master international macro modelers. You can give them your state-of-the-art DSGE model with all the bells and whistles, and they'll quickly X-ray it down to its most important bones. Why? Because they worked on papers like Cole and Obstfeld (1991 *JME*), Obstfeld and Rogoff (1995 *JPE*), Corsetti and Pesenti (2001 *QJE*), Devereux and Engel (2003 *ReStud*), and more. It is this highly tractable, very transparent work that—as massively simplified as it is relative to reality—helps us understand how key mechanisms work, and what we can expect when we embed those mechanisms in more complicated, more realistic models: Once I know how the basic mechanism works, if I add a piece to the machine and I see a difference in results, knowledge of the basics gives me a pretty good clue on where to start in order to understand the difference in results. If I start from the complicated machine right off the bat, it can be much more complicated to disentangle what is driving what.

Moreover, models are never supposed to be photographs of reality. They are supposed to be tools that help us understand key features of reality, important phenomena we want to focus on, things that we want to explain. If you give me a model that—thanks to the improved computing capabilities of our computers—does a fantastic job of reproducing all sorts of features of reality and even does a great job at forecasting, but I ultimately have no clue about the key driving forces of these amazing results, I similarly have no idea what to say to the student, colleague, central bank governor, or finance minister who's asking me “So, what explains these dynamics that we expect to see when policy action X is taken?” If you have built your fantastic model without actually working hard at understanding exactly how each piece contributes to its results, and without being able to explain it clearly to me and anyone else who's reading your work, we are not going very far. Yes, your fantastic model will reproduce all sorts of stuff and will be great at forecasting—for some time, maybe even a long time—but I will always be able to find some features of reality that will be of some relevance to some people that your model does not capture (as long as it is a model and not a perfect reproduction of reality itself, which I am not sure I would refer to as a model). And, sooner or later, human creativity will produce something that your model cannot forecast. A crisis

country Europe. In the EMU scenario, the European countries shared a common currency and central bank, but still had national, possibly uncoordinated fiscal policies. I showed how uncoordinated fiscal policies within EMU would have effects on the euro-dollar exchange rate, and I studied different scenarios for policy coordination across the Atlantic. I started working on this in 1991, long before there was any certainty that the euro would happen. The models in much of my work with Eichengreen and Giavazzi “descended” from the model in my thesis.

will happen, and bashers of all types will again be screaming that we have to throw away the entire toolkit and start back from square one.

So, to summarize, tractability for me is about working with tools that deliver results that I can understand and explain transparently in terms of the core mechanisms that are driving them. Mechanisms that can then be embedded as parts of more realistic, more empirically relevant models to be used, for instance, for policy analysis, with the hope of understanding results clearly and being able to explain them transparently to policymakers who do not have the time and do not want to hear about equations and the like.

I will conclude this note with Bob Tetlow’s reply to my tweet about Beatrice Cherrier’s first post on the tractability issue—where I was stating my hope/illusion to write my thoughts “in a few days.” You can see that I agree with Bob:

Fabio Ghironi @FabioGhironi · Apr 20
Interesting post by @Undercoverhist on the costs of pursuing tractability in econ models: beatricecherrier.wordpress.com/2018/04/20/wha... Hoping to write some thoughts on this in the next days.

What is the cost of 'tractable' economic models?
Edit: here's a follow-up post in which I clarify my definition of 'tractability' and my priors on this topic
Economists study cycles, but they also create some.
beatricecherrier.wordpress.com

3 replies 6 retweets 8 likes

Robert Tetlow @BobTetlow
Following

Replying to @FabioGhironi @Undercoverhist

Please do. On the surface of the issue, whatever the issue is, I'm at a loss to understand what the alternative is.

6:57 PM - 21 Apr 2018

2 Likes

2 likes

Update, 08/13/18

[Francisco Ilabaca asked a very important question](#) on August 10 after reading my note. Here's a picture for your convenience:



I finally made it to reply on August 12 by [tweeting a picture of my answer](#). I copy it below in case it can be useful to others:

This is a very important question. Based on my experience working with students, I think the point I was trying to make still applies. Many of them over the years went on the market with papers that contained quantitative models that had to be solved computationally. But the most successful strategy for the students was always to develop their models gradually, starting from the simplest versions, and adding ingredients, eventually converging to the setups that went into their JMPs. Those versions had more empirical, quantitative appeal than the simpler version(s) developed along the way, but having worked through the simpler versions allowed the students (and me!) to have a very clear understanding of how the various ingredients mattered for the results in the JMPs. When giving their presentations, the students were able to turn off/on model features to explain channel/mechanism clearly. Often, they were able to refer to analytical results that shed light on the workings of the JMP models. Understanding the models and the ability to convey this understanding built on both pencil-and-paper and computer work with successive versions of the models.

When you are on the job market, you are presenting to audiences that include people from many different fields. It is very important that you explain your results and the mechanisms that deliver them as clearly as possible even to people whose familiarity with macro may be limited to core courses they took as grad students long ago, and who have no familiarity at all with international economics. (This is less of a challenge if you are giving your seminar at, say, a Fed, where a much larger fraction of the audience will be field-insiders, and the need—at least occasionally in your seminar—to boil things down to the simplest mechanisms is much more likely to be less

pressing. However, the clarity of presentation from having gotten your hands dirty by going inside all mechanisms of your machine will always be valuable.)

Plenty of the background work from building your framework gradually won't go into your JMP draft, but it will help you nonetheless—at least, this was the experience with my students. It will help you understand what features of your model(s) are important to deliver its key qualitative implications, and what other features you may need to add to bring the model closer to data to perform quantitative exercises. If your paper includes an empirical part (increasingly important), the tighter you can make the connection between theoretical and empirical parts of the paper, the better. From my perspective, understanding all the nuts and bolts of your model facilitates that process too. In the end, the published version of your paper is likely to be very different from the version you will go on the market with: You will be making many adjustments based on the feedback you receive in seminars and from referees and editors. But experience suggests that my students who put much effort into gradual model-building were those whose papers were most clearly written and best received also in the publication process.

You are lucky to be doing your PhD in a department with an amazing group of macro and international economists. I hope my thoughts are helpful for you, but no one is in a better position to help you develop toward the job market (and beyond) than your advisors at UCI. It is a truly superb group. Enjoy the experience of learning from them!