Research Statement

Joseph B. Steinberg
Assistant Professor, Department of Economics, University of Toronto
September, 2019

I am a quantitative international macroeconomist. I use calibrated dynamic general equilibrium models to answer questions in open-economy macroeconomics and international trade such as: why do countries borrow and lend from one another, and how do international capital flows affect countries' domestic economies? How do changes in trade policy affect macroeconomic dynamics in the short and long run? Recently, I have also begun to use the tools of quantitative macroeconomics to study the effects of redistributive policies aimed at curbing wealth inequality. My research on these topics is unified by two goals: to provide evidence on pressing issues in modern policy debates; and to advance the methodological frontier by developing new models that more accurately depict the complexities of economic relationships between countries, between sectors, and between individual firms and households. In this research statement, I describe my contributions to each of these topics and my plans for future research.

1 Causes and effects of international integration

The international economic landscape has changed dramatically since the 1980s. International trade in goods and financial assets has burgeoned, complex global supply chains have emerged, and fast-growing emerging economies have saved heavily while slower-growing advanced economies have borrowed. Understanding the underlying drivers of these changes and their implications for labor markets, financial markets, and other arenas is of paramount importance for modern policymakers. My research in this area provides precise quantitative analysis of these issues using dynamic general equilibrium models with rich multi-sector, input-output production frameworks that accurately depict the process of globalization.

Since the early 1980s, the United States and other advanced economies have run large trade deficits and borrowed heavily from China and other emerging economies. These imbalances in trade and capital flows have become one of the most important empirical phenomena in the international economics literature and have featured heavily in broader public conversations about the consequences of globalization. One of the central questions surrounding these imbalances concerns their root cause. Some scholars argue that increased demand for saving in emerging economies—a global saving glut—is responsible, while others argue that reduced demand for saving in advanced economies—a domestic saving drought—is the culprit. In **"On the Source of U.S. Trade Deficits: Global Saving Glut or Domestic Saving Drought?" (*Review of Economic Dynamics*, 2019)**, I provide a quantitative resolution to this question by using my modeling framework to conduct a wedge accounting analysis of U.S. trade flows and macroeconomic dynamics. I find that increased demand for saving in the rest of the world explains virtually all of the cumulative U.S. trade deficit during the 1990s and 2000s.

Wedge accounting decomposes macroeconomic fluctuations into contributing factors using equilibrium models. Distortions, or wedges, that represent these factors are incorporated into a model's structure, values

are assigned to the wedges so that the model's equilibrium matches the data at hand, and then counterfactual equilibria are constructed in which the wedges are "turned on" one by one in order to isolate their contributions to the trends observed in the data. A number of previous studies have used small open economy models to account for the causes of trade imbalances, but the partial-equilibrium nature of these models prevents separate identification of the effects of domestic and foreign factors. By contrast, I use a general equilibrium accounting framework that features wedges in both the United States and the rest of the world that affect each region's demand for saving, which allows me to definitely determine whether domestic or foreign factors have caused U.S. trade deficits. My analysis also incorporates two exogenous driving forces: productivity growth and demographic change. Productivity has grown more rapidly in the rest of the world than in the United States and the population has grown relatively older, and, all else equal, these forces would have led the rest of the world to borrow heavily from the United States, rather than the other way around. My results indicate that saving wedges in the rest of the world explain 96 percent of the cumulative difference between observed U.S. trade deficits and the trade surpluses the U.S. would have run naturally as a result of productivity growth and demographic change between 1995 and 2011.

As the United States and other advanced economies have borrowed from the rest of the world, employment in manufacturing and other goods-producing sectors in these countries has fallen dramatically. Pundits, policymakers, and politicians often blame trade deficits for the decline in goods-sector employment, arguing that increased consumption of foreign goods has shifted these jobs offshore. The structural change literature offers an alternative explanation: rapid productivity growth has enabled goods-producing firms to produce their products with fewer workers as compared to other sectors. In **"Global Imbalances and Structural Change in the United States" (*Journal of Political Economy*, 2018)**, Timothy Kehoe, Kim Ruhl, and I use a multi-sector general equilibrium model to measure the extent to which each of these forces have reduced goods-sector employment in the United States. Contrary to popular wisdom, we find that trade deficits are responsible for only 15 percent of the decline in goods-sector employment since the early 1990s; instead, rapid productivity growth in the goods sector is the primary driver of this trend.

Our model combines an open-economy theory of intertemporal trade with a closed-economy theory of productivity-driven structural change, allowing us to run a kind of horse race between these two mechanisms. We calibrate our model's parameters to replicate the U.S. trade deficit and sectoral productivity growth rates during 1992–2012, but our model also matches the path of goods-sector employment during this period exactly even though we do not target this statistic, indicating that our model is well-suited to the task at hand. To measure the contribution of the trade deficit and asymmetric productivity growth to the decline in goods-sector employment, we construct counterfactual equilibria in which we turn off these two forces one at a time to isolate each of their effects. Our model and quantitative strategy have laid the groundwork for numerous studies on the relationships between international trade, sectoral reallocation, and macroeconomic dynamics, and our results, which imply that policies aimed at reducing the U.S. trade deficit are unlikely to bring jobs back to manufacturing and other goods-producing sectors, provide important context for the ongoing conversation about the consequences of globalization.

In addition to shedding light on the causes and effects of U.S. trade deficits, these two papers also make important contributions to our understanding about the relationship between these deficits and the substantial decline in real interest rates that has accompanied them. Conventional saving-glut logic dictates

that U.S. trade deficits and low interest rates are two sides of the same coin: as emerging economies' demand for saving rose, interest rates fell in order to clear the world saving market. In my 2018 *Journal of Political Economy* paper with Tim Kehoe and Kim Ruhl, we show that falling interest rates and trade deficits are largely unrelated; our model generates trade deficits without a corresponding decline in interest rates. In my 2019 *Review of Economic Dynamics* paper, I use a simple theoretical model to illustrate the explanation for this result: when domestic and foreign products are imperfect substitutes, the real exchange rate, not the real interest rate, adjusts to clear the market for saving. My quantitative analysis in this paper shows that increased demand for saving in the rest of the world explains virtually all of the dynamics of the U.S. real exchange rate in recent decades in addition to movements in the trade deficit, but that the decline in real interest rates is explained primarily by investment distortions.

International trade in assets has burgeoned in recent decades as well as trade in goods and services, but this financial globalization has had disparate effects across countries: advanced economies' wealth has grown more internationally diversified, but emerging economies' wealth remains heavily biased towards domestic assets. In **"International Portfolio Diversification and the Structure of Global Production" (*Review of Economic Dynamics*, 2018)**, I show that this asymmetry is a natural consequence of the changes in the global production structure that have occurred since the early 1990s. There are two main forces that drive this result. First, trade in goods and services makes foreign assets better hedges against domestic shocks, and so countries that become more open to trade also become less biased towards domestic assets. Second, larger countries, which account for larger shares of global asset market capitalization, naturally hold more domestic assets, and so faster-growing countries' portfolios become more biased towards domestic assets. For advanced economies, which have become more open to trade but have grown more slowly than the average country, both of these forces have reduced country portfolio home bias. For emerging economies, however, which have become more open to trade and have grown faster than average, these two forces have opposed one another, and consequently the diversification of these economies' country portfolios has remained stable over time.

To quantify the effects of these forces on country portfolios, I build a multi-country international business cycle model with an input-output structure that accurately captures international trade relationships as well as the distribution of output and absorption across countries. I calibrate this structure to match input-output matrices from different time periods to measure the overall effects of changes in this structure on country portfolios, and then I construct counterfactual matrices to isolate the effects of specific changes like trade openness and country size. To solve for equilibrium country portfolios, I develop a perturbation-based method that allows for arbitrary numbers of countries and assets. This method represents a significant improvement upon similar approaches used in other studies that work only for conventional two-country models.

2 Macroeconomic consequences of protectionist trade policies

After decades of deepening international economic integration, recent years have witnessed a resurgence of protectionism around the world: decades-old political unions and free trade agreements are under threat and trade wars have begun to drag down the global economy. My research in this area builds new quantitative

tools that synthesize insights from macroeconomics and international trade in order to analyze how rising trade barriers, or even just the possibility that trade barriers could rise in the future, affect macroeconomic dynamics.

The harbinger of the rise of protectionism was the so-called Brexit referendum in June, 2016, in which the United Kingdom voted to leave the European Union. The law that authorized this referendum was silent about the policies that will govern trade between the U.K. and the remainder of the E.U. in the future, creating considerable uncertainty about whether Brexit would be "hard," in which tariffs and other trade barriers rise dramatically because the U.K. leaves the European single market, or "soft," in which the U.K. retains single-market access and trade barriers change little. The Brexit vote was followed by widespread speculation that this uncertainty about future trade policies would cause immediate harm to the U.K. economy. In **"Brexit and the Macroeconomic Impact of Trade Policy Uncertainty" (*Journal of International Economics*, 2019)**, I use a dynamic model of trade between the U.K., the E.U., and the rest of the world to measure the macroeconomic cost of Brexit uncertainty. I find that although the overall cost of Brexit will be large, especially if hard Brexit occurs, the cost of uncertainty about Brexit is negligible.

In contrast to my findings about Brexit, empirical evidence suggests that trade policy uncertainty (TPU) can have significant macroeconomic effects. For example, several studies point to the surge in U.S. imports of Chinese goods that followed China's 2001 accession to the World Trade Organization; the United States had granted "most favored nation" status to China twenty years earlier on a temporary basis, which gave Chinese exporters access to U.S. markets at the same low tariff rates enjoyed by WTO members, and accession elimininated the possibility that this status could be revoked in the future. Definitively attributing changes in macroeconomic variables to uncertainty about trade policy is challenging, however, because it requires one to ascertain how these variables would have evolved in the absence of this uncertainty—in other words, it requires a theory of how TPU affects economic agents' behavior. The most common theory in the TPU literature hinges on a real option value mechanism: entering the export market requires large sunk costs, and when future trade policies are uncertain it is optimal for firms wait to incur these costs until this uncertainty is resolved.

In my analysis of Brexit, I embed a microeconomic model of export participation dynamics that captures this mechanism into a multi-country international macro model with stochastic trade costs that reflect the timing of the Brexit process. Model agents learn in 2016 that Brexit will occur three years later, but they do not learn until that time whether Brexit will be hard or soft. To measure the overall effect of Brexit, I compare the model's equilibrium to a counterfactual in which the Brexit vote does not occur at all, and to measure the effects of Brexit uncertainty, I compare the stochastic equilibrium to a pair of perfect-foresight counterfactuals—one each for hard and soft Brexit—in which agents learn immediately which form of Brexit they will face. I find that the overall welfare cost of Brexit will be substantial, ranging from 0.4 percent to 1.2 percent in consumption-equivalent terms, but that less than a quarter of a percent of this cost is attributable to uncertainty.

In addition to providing valuable evidence about the effects of Brexit, my approach to modeling trade adjustment dynamics and the effects of TPU significantly advances the methodological frontier. My theory of exporters' life cycles broadens and generalizes previous theories along several dimensions, and my model is the first to incorporate this kind of framework into a general equilibrium environment with many

asymmetric countries. Additionally, I develop a novel method to solve for my model's exact transition dynamics, rather than a linear approximation as is commonly employed in DSGE modeling, which allows me to accurately measure the effects of uncertainty.

In **"Comment on: 'The Economic Effects of Trade Policy Uncertainty' by Dario Caldara, Matteo Iacoviello, Patrick Molligo, Andrea Prestipino, and Andrea Raffo" (*Journal of Monetary Economics*, Forthcoming)**, I draw on the insights from my analysis of Brexit to comment on another recent paper in the TPU literature. The authors of this paper build a new firm-level dataset on TPU exposure, use this data to show that firms reduce investment when exposure to TPU increases, and use a sticky-price version of my Brexit model to show that nominal rigidities play a crucial role in transmitting the microeconomic effects of trade policy uncertainty to the macroeconomy. My comment on this paper makes two contributions. First, I link the authors' data to the U.S. input-output tables to glean some insights about the economic mechanisms underlying firms' concerns about TPU. My analysis indicates that the source of these concerns varies significantly across industries: firms in export-intensive industries worry that increased foreign trade barriers could reduce demand for their products; firms in import-competing industries hope that increased domestic trade barriers could reduce competition; and firms in industries that import intermediate inputs are concerned that increased domestic trade barriers could raise input prices. Second, I use a simple theory of exporters' price setting behavior to illustrate why nominal rigidities can exacerbate the effects of trade policy uncertainty in some circumstances, but can actually mitigate these effects in other circumstances.

Since the election of President Donald Trump, the United States has also taken a decisively protectionist turn, increasing tariffs on hundreds of billions of dollars of Chinese goods and threatening to terminate the North American Free Trade Agreement (NAFTA). In **"The Macroeconomic Impact of NAFTA Termination" (*Canadian Journal of Economics*, Forthcoming)**, I hone the tools I developed to analyze Brexit and use them to assess the short- and long-run effects that NAFTA termination would have on the U.S., Canadian, and Mexican economies. I find that regional trade flows would fall dramatically and that consumers in all three countries would lose, but that regional trade imbalances—the source of the Trump Administration's complaints about NAFTA—would remain.

Since NAFTA's inception, extensive regional supply chains have emerged, especially in the transportation equipment industry, and so accurately modeling the effects of NAFTA termination requires a realistic multi-sector, input-output production structure that can capture the complexity of these intersectoral and international relationships. Fortunately, my work on U.S. trade deficits and international portfolio diversification provides exactly this kind of framework. My model of NAFTA termination, which integrates this framework with the tools I developed to analyze trade dynamics in my work on Brexit, is the first of its kind in the literature. It features multiple sectors, multiple countries, trade in intermediate goods as well as final goods, heterogeneous firms that pay sunk costs to begin exporting, and adjustment frictions in trade and factor markets that slow intersectoral and international reallocations. I calibrate the model so that it matches the three NAFTA countries' current production and trade relationships, and then use it to trace out the transition dynamics that would occur if NAFTA were unexpectedly terminated.

In the long run, regional trade flows would fall dramatically and consumers in all three countries would be worse off; Canadian and Mexican consumers would bear the brunt of these losses, as their economies rely heavily on trade with the United States, while U.S. consumption would fall only slightly. It would take

many years for the region to transition to its new steady state, however, and modeling this transition as well as the the long-run impact is quantitatively important for measuring the overall welfare cost of NAFTA termination. Moreover, the forces that shape this transition actually play significant roles in determining the long-run effects as well as the short-run effects. For example, long-run welfare losses are almost ten percent lower in a version of the model in which firms can costlessly move into and out of the export market. These results highlight the importance of using dynamic models to quantify the effects of changes in trade policy. My model's dynamic structure also allows me to speak directly to the Trump administration's central complaint about NAFTA, which is that it has resulted in "unfair" trade imbalances between the United States and its neighbors. I find that NAFTA termination would cause the U.S. trade deficit with Canada to shrink, but it would also cause the deficit with Mexico to rise, and it would have little impact on the aggregate U.S. trade deficit.

3 Ongoing research and plans for the future

To date, my published research has focused on topics in international macroeconomics. As a graduate student, I was also interested in macroeconomic questions related to inequality and redistribution—Fabrizio Perri and I analyzed these issues in **"Inequality and Redistribution During the Great Recession" (Minneapolis Fed Economic Policy Papers, 2012)**—and I have recently begun to return to these issues in a new research agenda with Shahar Rotberg, one of my former PhD students.

Wealth inequality has risen dramatically in recent decades, and there are growing calls among progressive policymakers to tax rich households' wealth to curb this trend. U.S. Senator Elizabeth Warren, for example, has proposed a 2-percent tax on wealth above $50 million that has received extensive media coverage and the backing of several high-profile economists, and N.D.P. leader Jagmeet Singh has proposed a similar policy here in Canada. Proponents claim that, in addition to reducing wealth inequality, a progressive wealth tax would raise substantial revenues that could be transferred to low-income households, but the tax base from which these revenues are generated could shrink as inequality falls. Additionally, critics argue that a progressive wealth tax could have adverse macroeconomic consequences and that tax evasion could limit its efficacy.

In our new paper, **"Tax the Wealthy? A Dynamic General Equilibrium Analysis of Progressive Wealth Taxation,"** Shahar and I build a general equilibrium model with overlapping generations of heterogeneous households and imperfect wealth tax enforcement to tackle these issues. We calibrate our model to match the current distribution of wealth in the United States, and then use it to trace out the short- and long-run consequences of progressive wealth taxation. We find that a progressive wealth tax would substantially reduce wealth inequality, but it would raise much less revenue in the long run than in the short run, and it would also reduce capital, output, and wages. The vast majority of households would approve of a progressive wealth tax, but high-earning young entrepreneurs, not low-income workers, would reap the largest benefits. Our results also highlight another potential concern about wealth tax evasion: if rich households move some of their wealth offshore to evade the tax, the domestic capital stock will fall further, hastening and exacerbating the macroeconomic contraction, and leaving many workers worse off in the long run.

Shahar and I also plan to use our model to analyze optimal wealth redistribution policies. Our frame-

work allows us to determine the optimal level of progressivity—as captured by the threshold above which wealth is taxed—as well as the optimal marginal wealth tax rate, and to gauge the extent to which transition dynamics are important in shaping these results. I believe that this research agenda will yield important insights about the macroeconomic effects of wealth redistribution and will provide valuable context for ongoing policy debates.

In addition to my ongoing work on wealth redistribution, I continue to advance my research on trade adjustment dynamics. In preliminary work entitled **"Export Market Penetration Dynamics,"** I use microdata on Mexican exporters and the export participation dynamics model I developed in my work on Brexit to study how exporters' life cycles vary with the characteristics of the markets they serve. It is well-known in the trade literature that the size distribution of exporters varies systematically with export markets' characteristics: more firms export to larger, richer destinations than smaller, poorer ones, but sales are more concentrated amongst the largest exporters in the former than in the latter. In my empirical analysis, I use the Mexican microdata to document novel evidence that exporters' life cycles also vary systematically across destinations: new exporters are smaller and more likely to exit than incumbents in larger markets than in smaller ones. In my quantitative analysis, I show that my export participation dynamics model, which generalizes several of the most widely-used theories of exporting in the literature, accounts for both of these sets of facts.

The key ingredient in the model is the cost that a firm must pay to build its customer base in an export market, which depends on the number of new customers the firm attempts to reach, the firm's current customer base, and the market's overall size. In choosing how many foreign customers to acquire, firms weigh the marginal customer acquisition cost against the marginal profit earned by serving additional customers. The cost of reaching even a single customer is positive, which means that sufficiently unproductive firms decide not to export at all. Acquiring new customers is increasingly costly, which means that more productive firms have more customers, but acquiring customers is less costly when a firm has more customers to begin with, which means that firms gradually accumulate customers over time. It is also less costly to acquire customers in larger markets—there are returns to market size in advertising—so larger markets attract more exporters, each of which serves more customers than it serves in smaller markets. Despite its parsimony, this theory of export market penetration costs generalizes several of the most important theoretical contributions to both the static and dynamic international trade literatures.

The model has two sets of parameters: one that governs market penetration costs; and another that governs export destinations' characteristics (population, income, import tariffs, etc.). To demonstrate the model's ability to capture the patterns in the microdata, I calibrate the first set of parameters so that the model's stationary equilibrium replicates the behavior of Mexican firms that export to the United States, and then simulate the model for the other destinations in the dataset while holding the calibrated parameters fixed. The simulations exhibit the same systematic variation in the size distribution and life-cycle dynamics of exporters across destinations as in the data, indicating that the model accounts for both sets of facts.