Research Agenda: Ariel Burstein

November 2010

My research examines a variety of central questions in international macroeconomics and trade, such as: Does international trade improve welfare by stimulating innovation? To what extent can the increase in globalization account for the observed rise in wage inequality in developed and developing countries? What are the aggregate consequences of reallocating firm-embedded know-how across countries in the form of FDI? Why are relative prices across countries so volatile over time? A unifying theme of my research is to develop quantitative models and bring in new data to shed light on these classic questions. Here I take the opportunity to report on the progress of some of my recent and ongoing work in addressing some of these questions.

General Equilibrium Perspectives on Innovation by Firms

In this work, Andrew Atkeson and I explore a commonly held view that international trade has extra benefits because it stimulates innovation by firms. According to this idea, reductions in trade costs make large firms that export face a bigger global market, which increases their incentives to innovate to reduce costs. Hence, the larger exposure to trade leads to an increase in productivity by some firms, which contributes to raising aggregate productivity and welfare.

To assess this view, Andy and I (Atkeson and Burstein 2010a) develop a general equilibrium model that captures the dynamic decisions of heterogeneous firms to exit, export, and invest in innovation to both improve existing products (process innovation) and create new products (product innovation). This model can then be used to aggregate-up from firm-level decisions to obtain a deeper understanding of how aggregate productivity should be expected to respond in general equilibrium to changes in the economic environment such as international trade costs. Our model extends Hopenhayn's (1992) and Luttmer's (2007) model of firm dynamics with exit and entry of new firms to include a R&D decision by incumbent firms following Griliches' (1979) model of knowledge capital. We consider an open economy version of these models with fixed and marginal international trade costs, as in Melitz (2003).

We propose a simple algorithm to assess the impact of a change in the cost of international trade on aggregate productivity and welfare in the long run. A striking feature of our results is that, across a range of model specifications that we can solve analytically, consideration of the endogeneity of firms' decisions to exit, export, and invest in process innovation have no impact, to a first-order approximation, on our models' implications for the long-run impact in general equilibrium of a change in international trade costs on aggregate productivity. In fact, to a first-order approximation, one would obtain the same results for the long run change in aggregate productivity in a simple model in which firms' exit, export, and process innovation decisions were held fixed (completely inelastic) and only firms' entry decisions responded to changes in the economic environment.

This is not to say that there are no important changes in firm dynamics that result from a change in international trade costs. In fact, for certain assumptions about the elasticity of incumbent firms' investments in process innovation with regard to changes in the costs or benefits of such innovation, our

model predicts dramatic changes in these firms' innovative efforts and productivity growth rates. Cutting trade costs stimulates process innovation for exporting firms, as these firms seek to grow their profits by selling to a larger market at lower costs. Hence, a reduction in trade costs leads to a reallocation of production, export status, and investments in process innovation from smaller, less-productive, non-exporting firms to larger, more-productive, exporting firms and this reallocation leads to an increase in the productivity of the average firm and in the share of exports in output (some of these implications are further described in Burstein and Melitz 2010).

Yet, in general equilibrium, this reallocation does not have a first-order effect for the model's implications for long run aggregate productivity. The rise in the productivity of the average firm stemming from the improvements in process innovation and entry into export markets is offset by a reduction in product innovation. Why is that? If other firms in the economy become more productive as they innovate more, this drives down the profitability of entering firms that tend to be small non-exporters competing with globalized firms. This is the self-limiting nature of innovation: the stimulating effect of process innovation and entry into exporting is largely offset by a reduction in product innovation.

We evaluate the strength of this offset in a version of the model that relaxes some of the assumptions that give rise to our stark analytic result described above. In particular, we consider a parameterized version of our model that accounts for some salient features on the share of exporters in output and employment and the firm size distribution in the U.S. economy. In one of our experiments, we consider a reduction in trade costs that more than doubles the share of exports in output. This greatly stimulates the innovative activity of exporters, causing a surge in average productivity: long run average productivity rises by 7.5 times the percentage change in trade costs. But this increase in average productivity is almost entirely offset by a reduction in product innovation, which falls by a factor of 7.4 times the percentage change in trade costs. Moreover, this long run rise in aggregate productivity stemming from the increase in innovation intensity turns out not to make a big difference for welfare of the representative consumer. This is due to the fact that after firms increase their process innovation, it takes a long time to materialize and impact aggregate productivity.

Andy and I (Atkeson and Burstein 2010b) are currently using our framework to more generally assess the aggregate implications of changes in economic policies that affect the costs and benefits to firms of innovative activity. In order to allow a role for welfare-enhancing policies in our model, we include spillovers from R&D. Our model nests widely-used models of semi-endogenous and endogenous growth (see e.g. Acemoglu 2009) based on the value of a key parameter determining to what extent the productivity increase in some firms crowds-out profits of other firms. In the knife-edge case in which a firm's profits are independent of other firms' productivities, business R&D is the engine of endogenous growth. In contrast, if productivity increase in incumbent firms crowds-out profits of other firms, business R&D is not the engine of growth, as in semi-endogenous growth models.

When business R&D is not the engine of growth, we show that the calculation of the impact of a wide array of policies directed at inducing innovation by firms on aggregate productivity in the long run can be boiled down to a relatively simple accounting procedure that is straightforward to implement.

Namely, to a first-order approximation, the aggregate effects of a subsidy on innovation are the same as the effects of a subsidy on variable profits, keeping innovation decisions of incumbent firms unchanged. This procedure can be understood as follows. First, subsidies on innovation are equivalent to subsidies on firms' variable profits because subsidizing the return to innovation is equivalent to subsidizing the cost of innovation. Second, any increase in innovation by incumbents lowers profits of entrants, reducing product innovation and keeping aggregate productivity unchanged. Once again, macroeconomic forces limit the aggregate impact of a change in incumbents' innovation decisions.

We identify the key parameters in our model that shape the effects of innovation policies, and revisit previous empirical approaches to estimating these effects through the lens of our model. As suggested above, data on the response of R&D of incumbent firms to changes in innovation subsidies is not informative on the aggregate effects of innovation policies. Moreover, because the crowding out effects we identify arise in general equilibrium, data on the response of output of individual industries to changes in innovation policies is informative only as long as the comprehensiveness of these policy changes across industries is known. This is because, in our model, the effects on output in an industry from an innovation subsidy that affects only this industry are very different from those of a subsidy to a large number of industries.

Finally we ask, as does Griliches (1998), whether innovation policies and innovation spillovers can generate large differences in per capita income and measured TFP across countries and time. As discussed in Prescott (1997) and Klenow and Rodriguez-Clare (2005), in constructing a theory of innovation or other unmeasured investment that can generate large differences in income per capita across countries, two important challenges for such a theory are that (1) it does not require implausibly large differences in innovation or unmeasured capital investment rates across time or countries, and (2) it does not require implausibly long transition dynamics from one steady-state to another.

Globalization and the skill premium

To what extent can the growth of trade and multinational production (MP) account for the rise in the skill premium in developed and developing countries? What are the different implications for the skill premium in developed countries of globalization with developing countries versus globalization with developed countries?

The classic model to study the link between international trade and inequality is the Heckscher-Ohlin (H-O) model. According to this theory, trade liberalizations shift factors of production between sectors towards a country's comparative advantage sector and raise the relative return to the factor that is used intensively in this sector --- the *between effect*. International trade increases the skill premium in countries that have a comparative advantage in skill-intensive sectors, and lowers the skill premium in countries that have a comparative advantage in unskill-intensive sectors. Previous empirical work (see e.g. Goldberg and Pavnick 2007), however, has cast doubt on the importance of trade in accounting for the skill premium because (i) inequality has increased both in countries abundant in unskilled labor and in countries abundant in skilled labor, and (ii) most factor reallocation occurs within rather than across sectors. An alternative interpretation of this evidence, however, is that the standard H-O model

abstracts from other potentially important channels beyond the between effect through which globalization affects the skill premium.

To allow for such additional channels, Jonathan Vogel and I (Burstein and Vogel 2010a) extend the H-O model in three dimensions. First, motivated by the large heterogeneity in size and export status within sectors, we introduce productivity heterogeneity across producers within sectors. Second, motivated by the observation that, within sectors, exporters tend to be more skill-intensive than non-exporters, we allow for skill-biased technology at the producer level. Third, we allow for producers to use their technologies to produce in foreign countries, engaging in multinational production (MP). MP is an important form by which producers supply foreign countries beyond international trade (for example, sales of U.S. foreign affiliates are more than twice as large as the value of U.S. exports).

Following a reduction in trade and MP costs, our extended model predicts within sector factor reallocation from less productive non-exporters to more productive exporters, as in recent international trade models with producer heterogeneity. If a producer's productivity is positively correlated with its skill intensity, then this reallocation of labor across producers within sectors raises the demand for skilled labor and the skill premium --- the *within effect*. In contrast to the standard Hecksher-Ohlin model, the within effect implies that the skill premium may rise in all countries. Hence, our extended model gives globalization (both trade and MP) a better chance to account for the rise in the skill premium in developed and developing countries. Which force dominates and by how much is a quantitative question that we address in our quantitative analysis.

We consider a parameterized four-country version of our model (matching U.S. trade and MP volumes in 2006 with skill-abundant and scarce countries) to conduct a series of counterfactuals. To isolate the role of international trade in shaping the skill premium, we first consider a reduction in trade costs moving from autarky to the level of trade in 2006, holding all other exogenous variables fixed (including factor endowments) and abstracting from MP. This is a "but for" analysis: What would the skill premium be, but for the availability of international trade opportunities? The rise in the skill premium caused by trade is 1.8% in the U.S. and 2.9% in skill-scarce countries.

The two central messages from these results are as follows. First, in contrast to the standard H-O model, our parameterized model is consistent with a rising skill premium in all countries. The skill premium rises in all countries because the within effect is relatively strong compared to the between effect. The between effect is weak because, in our parameterization as in the data, the factor content of trade in the U.S. is not very high (in fact, under some assumptions of our model, the between effect is fully determined by the factor content of trade --- see also Burstein and Vogel 2010b). Given that the between effect is weaker than the within effect, how much the U.S. trades matters more for its skill premium than with whom the U.S. trades (i.e. whether it trades with Europe or China). Second, the magnitudes of the changes in the skill premium of moving from autarky to 2006 levels of trade are quite small relative to, for example, the 24% rise in the (composition-adjusted) U.S. college-high school wage gap between 1966 and 2006 (see, e.g., Acemoglu and Autor 2010).

The relatively small trade share in the U.S. plays a critical role in explaining the relatively small impact of trade on the U.S. skill premium in our model. The share of MP in output, however, is at least twice as high. When we simulate a reduction in both trade and MP costs moving from autarky to the levels of trade and MP in 2006, the rise in the skill premium is much larger: 4.8% in the U.S. and 6.5% in the skill-scarce countries. MP is at least as important as international trade for determining the impact of globalization on the skill premium.

In order to assess the extent to which the growth of trade and MP can account for the rise in the skill premium between 1966 and 2006 in the U.S., we consider a second counterfactual in which we choose parameters to match the growth of trade and MP between these years. In this counterfactual we do not hold endowments or technologies fixed, but instead we target the increase in the supply of skilled labor and the greater growth of the skill-scarce countries between 1966 and 2006 and we allow for exogenous skill-biased technology growth to match the 24% increase in the U.S. skill premium. In our preferred, baseline parameterization, we show that in the absence of globalization, the rise in the skill premium in the U.S. would have been 1/9th smaller than the observed rise in the skill premium over this time period (in a less conservative parameterization in which we choose parameter values so that, given trade and MP shares, the between and within effects are strengthened, we find that in the absence of globalization the rise in the skill premium in the US would have been 1/5th smaller than the observed rise over this time period).

Whereas we use a structural, parameterized model to quantify the impact of international trade and MP on the skill premium, the empirical literature has mostly focused on three alternative approaches that emphasize: the factor content of trade, the extent of between-sector factor reallocation, and the mandated wage equation. We show that while each of these alternative approaches may provide estimates of the impact of international trade on the skill premium via the between effect, they do not capture the impact of the within effect of trade and MP. Using data generated by our model, in which the within effect is relatively strong, we show that these approaches underestimate the rise in the skill premium stemming from the increase in globalization.

In this work we capture two important forces in the debate on globalization and the skill premium---the between and within effects---and incorporate both trade and MP. In work with Javier Cravino (Burstein, Cravino, and Vogel 2010), we study an additional channel by which international trade can generate a rise in the skill premium. Capital equipment, such as computers and industrial machinery, embodies skill-biased technology and, hence, is complementary to skilled labor. Many countries, by importing a large share of their capital, import skill-biased technology and a rise in the skill premium. Our current work abstracts from the endogenous response of factor-specific innovation, which is potentially an important consideration to more fully assess the quantitative effects of globalization on inequality.

Understanding fluctuations in international relative prices

Why are relative prices across countries, as measured by real exchange rates (RERs), so volatile over time, and why do they so closely track movements in nominal exchange rates? These questions are at

the heart of the discussions on optimal exchange rate policy (i.e.: by depreciating the exchange rate, countries may be able to reduce the relative price and raise the output of the goods they produce) and on the role of the border in creating frictions to the international trade of goods (i.e.: if trade costs were small, international trade should arbitrage large differences in relative prices across locations). The high correlation between RER and nominal exchange rate movements has been widely suggestive of nominal macroeconomic non-neutralities (see e.g. Mussa 1986).

As a starting point to understand these movements in international relative prices, consider simple models of price setting with perfect competition or imperfect competition and constant markups, in which prices change one-to-one with movements in marginal production costs. According to these models, the relative price of a traded good produced in a common location and sold in two countries (subject to a linear transportation technology) should remain constant over time. This is the hypothesis of relative purchasing power parity (relative PPP). However, a large body of empirical work suggests that relative PPP does not provide an accurate representation for movements in relative prices of many goods (see e.g. the survey in Goldberg and Knetter 1995). Another implication of models with constant markups is that, if all goods can be traded at no cost across countries, consumer-price-based RERs should be constant over time. However, RERs are very volatile, even for tradeable goods (see e.g. Engel 1999).

In order to account for the observed deviations of relative PPP and the large movements of RERs for tradeable goods, researchers have considered two key sources of departure from the constant markup costless-trade benchmark. First, many goods are not traded (see e.g. Burstein, Eichenbaum, and Rebelo 2005), and even traded goods include a substantial non-traded distribution component (see e.g. Burstein, Neves, and Rebelo 2003). Changes in relative prices across countries for these goods may simply reflect movements in relative production or distribution costs across locations. Second, movements in international relative prices may be explained by the decision of individual firms to engage in pricing-to-market --- that is, to systematically vary over time the markup at which they sell their output in different locations (see e.g. Dornbusch 1987 and Krugman 1987). These movements in relative markups can stem from nominal rigidities that leave prices in the buyer's currency unchanged to exchange rate movements, or from some feature of demand or competition that gives rise to movements in demand elasticities (and hence markups) across destinations.

Identifying whether movements in relative prices across countries reflect movements in relative markups or in relative costs is a key ingredient for the design of optimal exchange rate policy. Movements in relative prices of individual goods across countries that reflect changes in relative costs across locations are typically efficient, while movements in relative prices that reflect changes in relative markups across locations are typically inefficient (see e.g. Engel 2010). Hence, changes in exchange rates have very different normative implications in both models.

Nir Jaimovich and I (Burstein and Jaimovich 2009) take the challenge of documenting the role of pricing-to-market in accounting for the observed movements in international relative prices. We do so using detailed information on prices in Canada and the U.S. at the level of individual products. In particular, we use scanner data for the period 2004-2006 from a major retailer that sells primarily

nondurable goods in multiple locations in Canada and the U.S. For each product, we observe the retailer's purchase cost from the vendor, i.e. the wholesale price, in each location and over time. We also identify the country of production of individual products that are sold in Canada and the U.S. Under the assumption that goods produced in a common location and sold in multiple locations are subject to common percentage changes in the marginal cost, movements in relative prices across locations for these goods must arise from changes in relative markups. With this information, we can thus assess the extent to which movements in relative prices of individual products across locations reflect the practice of pricing-to-market by producers and wholesalers.

Our findings demonstrate that pricing-to-market plays an important role in accounting for movements in international relative prices. We show that movements in aggregate RERs, constructed by averaging changes in relative prices across countries (expressed in a common currency) over a large set of products sold in both Canada and the U.S., closely track the large rise in Canada-U.S. relative unit labor costs over our sample period (which are mainly accounted for by the appreciation of the Canadian dollar against the U.S. dollar). For nontraded goods that are produced in each country and sold in both countries, these movements in aggregate RERs can simply reflect changes in relative costs across countries. However, the fact that this pattern holds as well for traded goods produced in a common location and sold in both countries implies that in response to the appreciation of Canada-U.S. labor costs, markups in Canada increase systematically relative to markups in the U.S.

Pricing-to-market does not stem, in a pure accounting sense, from large movements in nominal exchange rates and small movements in nominal prices in each country. Instead, nominal prices of individual products change frequently and by large magnitudes. Moreover, changes in international relative prices at the level of individual products, product-level RERs, are very large, roughly four times as volatile (at quarterly frequencies, excluding temporary price changes in each location) as the Canada-U.S. nominal exchange rate, even for traded goods. Hence, while cross-country differences in markups on average track movements in nominal exchange rates and relative labor costs, the idiosyncratic product-specific component of pricing-to-market is significant. We also show that the idiosyncratic and aggregate components of pricing-to-market are much more prevalent across countries than across regions within countries.

Our empirical findings relate to a recent and rapidly growing literature documenting the behavior of international relative prices for tradeable (but not necessarily traded) goods using detailed product-level information (see e.g. Crucini and Shintani 2008, Broda and Weinstein 2007, and Gopinath, Gourinchas, Hsieh, and Li 2010). Our empirical contribution is to measure the extent to which movements in relative prices of matched individual products across locations reflect pricing-to-market by producers and wholesalers, which can do because of two unique features of our data. First, by observing wholesale prices, we can more accurately measure movements in relative markups at the producer level than if we used retail prices, which contain a significant non-traded distribution component. Second, by using information on the country of production of individual products, we can identify goods that are actually traded, and infer changes in relative markups for these goods from observed movements in relative prices across locations. Fitzgerald and Haller (2008), like us, also find evidence of pricing-to-market by exporters in response to exchange rate movements using detailed firm-level data.

Our empirical findings raise some important questions: Why do relative markups systematically track movements in relative costs across countries, even if nominal prices of individual products change frequently and by large amounts? Why is pricing-to-market more prevalent across countries than within countries? Nir and I address these questions using a simple model of pricing-to-market and international trade with flexible prices (that extends my previous work, Atkeson and Burstein 2007, 2008). The international border plays an important role in our model by segmenting competitors across countries (more so than within countries), leading to the practice of pricing-to-market by exporters that optimally change their markup across locations (more so across countries than within countries) in response to idiosyncratic shocks and changes in aggregate relative labor costs.

Our model is highly stylized in order to gain analytical tractability and to help us identify key forces that can account for the observed movements in product-level and aggregate RERs. In doing so, we abstract from important industrial organization considerations (for models of incomplete pass-through that incorporate some of these considerations, see e.g. Goldberg and Hellerstein 2007 and Nakamura and Zerom 2010). An important question for future research is whether richer models of demand and market structure can give rise to idiosyncratic and aggregate movements in relative markups across locations (i.e. differences in pass-through across locations) like the ones observed in the data.

References

Acemoglu, Daron. 2009. “Introduction to Modern Economic Growth.” Princeton University Press, Princeton, NJ.

Acemoglu, Daron, and David Autor. 2010. "Skills, Tasks and Technologies: Implications for Employment and Earnings." Handbook of Labor Economics Volume 4, forthcoming.

Atkeson, Andrew and Ariel Burstein. 2007. "Pricing-to-Market in a Ricardian Model of International Trade." American Economic Review, 97(2): 362-367.

Atkeson, Andrew and Ariel Burstein. 2008. "Pricing-to-Market, Trade Costs, and International Relative Prices." American Economic Review, 98(5): 1998-2031.

Atkeson, Andrew, and Ariel Burstein. 2010a. “Innovation, Firm Dynamics, and International Trade.” Journal of Political Economy, 118(3): 433:484

Atkeson, Andrew, and Ariel Burstein. 2010b. “Aggregate Implications of Innovative Policy” Working paper. UCLA

Broda, Christian, and David Weinstein. 2007. "Understanding International Price Differences using Barcode Data." Mimeo. Columbia University.

Burstein, Ariel, Javier Cravino, and Jonathan Vogel. 2010. "Importing Skilled-Biased Technical Change." Working Paper. UCLA.

Burstein, Ariel, Joao Neves, and Sergio Rebelo. 2002. "Distribution Costs and Real Exchange Rate Dynamics During Exchange-Rate-Based Stabilizations." Journal of Monetary Economics, 50(6): 1189-1214.

Burstein, Ariel, Martin Eichenbaum, and Sergio Rebelo. 2005. "Large Devaluations and the Real Exchange Rate." Journal of Political Economy, 113(4): 742-784.

Burstein, Ariel and Nir Jaimovich. 2009. "Understanding Movements in Aggregate and Product-Level Real Exchange Rates." Mimeo. UCLA.

Burstein, Ariel and Marc Melitz. 2010. “Trade Liberalizations and Firm Dynamics.” Working Paper. UCLA.

Burstein, Ariel and Jonathan Vogel. 2010a. "Globalization, Technology, and the Skill Premium: A Quantitative Analysis." NBER Working Paper 16549. National Bureau of Economic Research.

Burstein, Ariel and Jonathan Vogel. 2010b. "International Trade Patterns, the Skill Premium, and Heterogeneous Firms." Mimeo. Columbia University.

Crucini, Mario and Mototsugu Shintani. 2008. "Persistence in law of one price deviations: Evidence from micro-data." Journal of Monetary Economics, 55(3): 629-644.

Dornbusch, Rudiger. 1987. "Exchange Rates and Prices." American Economic Review, 77(1): 93-106.

Engel, Charles. 1999. "Accounting for U.S. Real Exchange Rate Changes." Journal of Political Economy, 107(3): 507-538.

Engel, Charles. 2011. “Currency Misalignments and Optimal Monetary Policy: A Reexamination.” American Economic Review, forthcoming.

Fitzgerald, Doireann and Stefanie Heller. 2008. "Exchange Rates and Producer Prices: Evidence from Micro-Data." Mimeo. Stanford University.

Goldberg, Pinelopi K., and Rebecca Hellerstein. 2007. "A Framework for Identifying the Sources of Local Currency Price Stability with an Empirical Application." NBER Working paper 13183.

Goldberg, Pinelopi K., and Michael Knetter. 1997. "Goods Prices and Exchange Rates: What Have We Learned?" Journal of Economic Literature, 35(3): 1243-1272.

Goldberg, Pinelopi Koujianou and Nina Pavcnik. 2007. "Distributional Effects of Globalization in Developing countries." Journal of Economic Literature, 45(1): 39-82.

Gopinath, Gita, Pierre-Olivier Gourinchas, Chang-Tai Hsieh, and Nicholas Li. 2010. “International Prices, Costs and Markup Differences.” American Economic Review, forthcoming.

Griliches, Zvi. 1979. "Issues in Assessing the Contribution of Research and Development to Productivity Growth." Bell Journal of Economics, 10(1): 92-116.

Griliches, Zvi. 1998. "The Search for R&D Spillovers." NBER Chapters, in: R&D and Productivity: The Econometric Evidence. National Bureau of Economic Research

Hopenhayn, Hugo. 1992. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." Econometrica 60(5): 1127-1150.

Klenow, Peter J. and Andres Rodriguez-Clare. 2005. "Externalities and Growth." Handbook of Economic Growth, in: Philippe Aghion & Steven Durlauf (ed.), Handbook of Economic Growth, edition 1, volume 1, chapter 11, pages 817-861.

Krugman, Paul. 1987. "Pricing to Market When the Exchange Rate Changes." in Real Financial Linkages Among Open Economies, eds. S. W. Arndt and J. Richardson, London: MIT Press.

Luttmer, Erzo G.J. 2007. "Selection, Growth, and the Size Distribution of Firms." Quarterly Journal of Economics, 122(3): 1103--1144.

Melitz, Marc. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity" Econometrica. 71(6): 1695-1725.

Mussa, Michael. 1986. "Nominal Exchange Rate Regimes and the Behavior of Real Exchange Rates." Carnegie-Rochester Series on Public Policy, 25: 117-214.

Nakamura, Emi and Dawit Zerom. 2010. "Accounting for Incomplete Pass-through." Review of Economic Studies. 77(3): 1192-1230.

Prescott, Edward. 1997. "Needed: a theory of total factor productivity." Staff Report 242, Federal Reserve Bank of Minneapolis.