

THE INCREASING RETURNS REVOLUTION IN TRADE AND GEOGRAPHY

Prize Lecture, December 8, 2008

by

PAUL KRUGMAN

Princeton University, Woodrow Wilson School, Princeton, NJ 08544-1013, USA.

Thirty years have passed since a small group of theorists began applying concepts and tools from industrial organization to the analysis of international trade. The new models of trade that emerged from that work didn't supplant traditional trade theory so much as supplement it, creating an integrated view that made sense of aspects of world trade that had previously posed major puzzles. The "new trade theory" – an unfortunate phrase, now quite often referred to as "the old new trade theory" – also helped build a bridge between the analysis of trade *between* countries and the location of production *within* countries.

In this paper I will try to retrace the steps and, perhaps even more important, the state of mind that made this intellectual transformation possible. At the end I'll also ask about the relevance of those once-revolutionary insights in a world economy that, as I'll explain, is arguably more classical now than it was when the revolution in trade theory began.

1. TRADE PUZZLES

In my first year as an assistant professor, I remember telling colleagues that I was working on international trade theory – and being asked why on earth I would want to do that. "Trade is such a monolithic field," one told me. "It's a finished structure, with nothing interesting left to do."

Yet even before the arrival of new models, there was an undercurrent of dissatisfaction with conventional trade theory. I used to think of the propagation of this dissatisfaction as the trade counterculture. There were even some underground classics. In particular, Staffan Burenstam Linder's *An Essay on Trade and Transformation* (1961), with its argument that exports tend to reflect the characteristics of the home market, was passed hand to hand by graduate students as if it were a *samizdat* pamphlet. And there was also an important empirical literature on intra-industry trade, notably the work of Balassa (1966) and Grubel and Lloyd (1975), that cried out for a theoretical framework.

Why did the trade counterculture flourish despite the apparent completeness of conventional trade theory? Call it the similar-similar problem: the

huge role in world trade played by exchanges of similar products between similar countries, exemplified by the massive two-way trade in automotive products between the United States and Canada.

In 1980, this similar-similar trade was still a relatively new phenomenon. Trade in the first great age of globalization – the age made possible by steam engines and telegraphs – was mainly dissimilar-dissimilar: trade in dissimilar goods between dissimilar countries. Comparative advantage, which one may define as the idea that countries trade to take advantage of their differences, clearly explained most of what was going on. It was only with the recovery of trade after World War II, and especially after the major trade agreements of the 1950s and 1960s, that the more puzzling trade patterns that fed the counterculture became prominent.

Figures 1, 2, and 3 make this point, using data on British trade. Figure 1 shows the commodity composition of British exports and imports on the eve of World War I. The pattern of trade made perfect sense in terms of classical comparative advantage: Britain, a densely populated nation with abundant capital but scarce land, exported manufactured goods and imported raw materials.

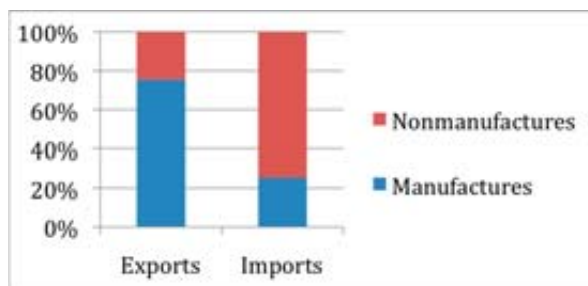


Figure 1. Composition of British trade circa 1910. Source: Baldwin and Martin, 1999.

By contrast, Figure 2 – which shows comparable data for 1990 – offers no comparably easy interpretation: Britain both imported and exported mainly manufactured goods. One might have hoped that a look within the manufacturing sector would reveal a clearer pattern, but this brings us to the issue of intra-industry trade: trade in manufactured goods, especially between countries at similar levels of development, consists to a large extent of two-way exchanges within even narrowly defined product categories.

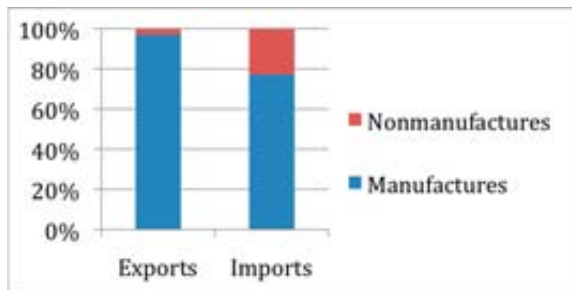


Figure 2. Composition of British trade in the 1990s. Source: Baldwin and Martin, 1999.

And trade, as reconstituted after World War II, *did* take place to a large extent between similar countries, much more so than in the first age of globalization. Figure 3 illustrates this crudely, comparing Britain’s trade with Europe and with rest of the world in 1913 and 1990.




Figure 3. Destination of British exports. Source: Baldwin and Martin, 1999.

Before World War I, Britain traded remarkably little, by modern standards, with its neighbors, instead focusing on distant lands able to produce what Britain could not – cheap wheat and meat, tea, jute, and so on. By 1990, however, while such trade had by no means vanished, Britain was part of a European economy in which nations seemingly made a living by taking in each others’ washing, buying goods that they could and, at least as far as the statistics indicated, did produce for themselves.

So what was going on?

2. INCREASING RETURNS AND TRADE

It shouldn’t have been that hard to make sense of similar-similar trade. Indeed, for some economists it wasn’t. In his seminal paper on the rise of intra-industry trade in Europe, Balassa (1966) stated it quite clearly: each country produced only part of the range of potential products within each industry, importing those goods it did not produce, because “specialization in narrower ranges of machinery and intermediate products will permit the

exploitation of economies of scale through the lengthening of production runs.”

Yet this straightforward-seeming explanation of similar-similar trade was not at all part of the standard corpus of international trade theory circa 1975. It was not so much that these ideas were rejected as that they seemed incomprehensible. Why?

The answer was that unexhausted economies of scale at the firm level necessarily imply imperfect competition, and there were no readily usable models of imperfect competition to hand. Even more to the point, there were no *general equilibrium* models of imperfect competition readily to hand – and trade theory, perhaps more than any other applied field of economics, is built around general equilibrium analysis.

The result was the state of affairs almost triumphantly described by Harry Johnson (1967): “The theory of monopolistic competition has had virtually no impact on the theory of international trade.”

Then came the new monopolistic competition models: Lancaster (1979), Spence (1976), and above all Dixit and Stiglitz (1977). All of these papers were intended by their authors as ways to address the classic welfare questions about whether monopolistic competition led to inefficient scale, or perhaps to production of the wrong mix of products. But when I learned about the new literature (in a short course taught by Robert Solow in 1976), I – like a number of other people working independently, including Victor Norman (1976) and Kelvin Lancaster (1980) himself – quickly saw that the new models provided “gadgets”, ways to think about the role of increasing returns in a variety of contexts. And there was, in particular, a near-perfect match between simple models of monopolistic competition and the stories already circulating about intra-industry trade.

It quickly became apparent (Norman, 1976; Krugman, 1979; Lancaster, 1980) that one could use monopolistic competition models to offer a picture of international trade that completely bypassed conventional arguments based on comparative advantage. In this picture, countries that were identical in resources and technology would nonetheless specialize in producing different products, giving rise to trade as consumers sought variety. A natural extension – although, like many things that seem obvious in retrospect, it was surprisingly hard at first to figure out how to do it – was to bring comparative advantage back in. This was most easily done by assuming that all the differentiated products within an industry were produced with the same factor proportions; one could then explain *inter*-industry specialization in terms of Heckscher-Ohlin, with an overlay of *intra*-industry specialization due to increasing returns. And this extension, as represented for example by Helpman (1981) and Dixit-Norman (1980), in turn meant that the new models offered an intellectually satisfying explanation of similar-similar trade: similar countries had little comparative advantage with respect to each other, so their trade was dominated by intra-industry trade caused by economies of scale.

What was really needed to get peoples' attention, however, was a "killer app": a demonstration that the new view offered a fundamentally different insight into something that mattered. I found that killer app in an empirical insight by Balassa (1966), who pointed out that trade liberalization among industrial countries had proved surprisingly non-disruptive, belying fears that, for example, there would be a major rearrangement of Europe's industrial landscape after the formation of the Common Market, and possibly large effects on income distribution. Because trade expansion had taken the form of intra-industry specialization rather than inter-industry specialization, Balassa noted, "the fears expressed in various member countries of the Common Market concerning the demise of particular industries have not been realized. There are no examples of declining manufacturing industries in any of the member countries."

In Krugman (1981), I built a special version of the emerging style of model to encapsulate this phenomenon. What the model showed was that the classic Stolper-Samuelson result, in which trade liberalization hurts scarce factors, *can* emerge – but only if comparative advantage is strong and/or economies of scale weak. In the reverse case, which seemed to describe the growth of trade among industrial countries, trade was win-win.

There was one more significant insight from the application of Dixit-Stiglitz-based models to trade: Burenstam Linder was right! Once one added transport costs to the model, it was straightforward to show that countries would, other things equal, tend to become exporters in the industries in which they had large domestic markets. As is so often the case, the logic of this result was obvious once the result had been devised, but not at all obvious beforehand. Indeed, I began the research that led to Krugman (1980) with the strong presumption that countries would *not* tend to export goods for which they had a large home market – I came to bury Burenstam Linder, not to praise him. But the algebra said otherwise, and the intuition followed. Increasing returns provide an incentive to concentrate production of any one product in a single location; given this incentive to concentrate, transport costs are minimized by choosing a location close to the largest market, and this location then exports to other markets.

Initially, the "new trade theory" seemed to consist of a series of special-purpose, incompatible models. It turned out, however, that it was possible to create a common ground for many though not all of the models, and extend that common ground to much traditional trade theory as well, using an insight originally due to Paul Samuelson (1949). In explaining factor price equalization, Samuelson reversed the usual way we think about trade, as a process of coming together. Instead, Samuelson thought of trade as the result of a process of coming apart. He envisaged a Tower of Babel scenario, in which an angel descends from Heaven and breaks up a previously unified economy: factors of production suddenly find themselves with national labels, and are only able to work with other factors that have the same national label. Samuelson pointed out that factor price equalization would take place if and only if the international distribution of factors of production was such

that it was possible, even while obeying the angel's new limits, to reproduce the production of the pre-angel integrated economy – and that in such circumstances, specialization and trade could be viewed as being *about* reproducing the integrated economy.

Elhanan Helpman and I (1985) used the same approach to think about trade involving both comparative advantage and increasing returns. The key insight was that in order to reproduce the integrated economy, it was necessary to locate all production of each good subject to economies of scale within one country. This approach united factor-proportions-based comparative advantage and specialization due to economies of scale: both could be viewed as part of how a world economy cursed by Samuelson's angel undid the damage.

This approach also, more or less *en passant*, made it clear that increasing returns ordinarily reinforce, rather than call into question, the argument that there are gains from trade. To be sure, in cases where the integrated economy is not reproduced, it's possible to conjure up examples in which countries are worse off with trade than without, in a way that isn't possible in pure comparative advantage models. But the clear presumption is that trade is a good thing under increasing returns – indeed, better than previously thought.

By the mid 1980s, then, the “new trade theory” had integrated increasing returns more or less seamlessly into our understanding of international trade. The impossible complexity that had previously daunted economists contemplating a major revision of trade theory had vanished, replaced by a surprisingly simple and elegant structure.

But how did that happen? Why did the problems facing the trade counterculture seem to melt away? I'd argue that at the heart of the story was an attitude shift on the part of international economists.

3. SOME META REFLECTIONS

The emergence of the new trade theory was, in the first place, made possible by the new models of monopolistic competition. But it did not remain confined to those models; by the mid-1980s recognizably “new trade” approaches had been taken to trade involving external economies, Cournot and Bertrand oligopoly, even contestable markets. What made it all possible was a shift in attitude among trade theorists, mainly consisting of two changes. First, there was a new willingness to explore the implications of illuminating special cases rather than trying to prove general results given some broad upfront assumptions. Second, there was a change in focus from detailed predictions – which country produces each specific good – to system-level or aggregate descriptions of the pattern of trade.

On the first point: in the late 1970s many trade theorists thought of themselves as theorem-provers. Given big initial assumptions such as constant returns, so many factors, etc., what could be proved true about trade, specialization, and welfare? There was, at least in the theoretical literature, a

push for generality. Yet this generality was mainly spurious: the results were general *given* the big assumptions, but those assumptions were, in fact, highly restrictive, ruling out much of what was obviously true about real-world trade and specialization.

The new trade theory, instead, focused on strongly special, even silly-seeming cases. (“Dare to be silly” became one of my principles for research.) There is no good reason to believe that the assumptions of the Dixit-Stiglitz model – a continuum of goods that enter symmetrically into demand, with the same cost functions, and with the elasticity of substitution between any two goods both constant and the same for any pair you choose – are remotely true in reality. The assumptions are instead chosen, with full self-consciousness, to produce a tractable example that contains what older trade theories left out – namely, the possibility for intra-industry specialization due to economies of scale.

The use of deliberately unrealistic assumptions is, of course, common in much of economics. Nonetheless, I can report from early experience that the new style of modeling was met with considerable hostility at first. Some discussants dismissed the whole enterprise as obviously pointless, given the unrealism of the setup. Some even insisted that if all the goods enter utility the same way, they must be perfect substitutes. There was a widespread sense that the new trade theorists were cheating.

Meanwhile, even with all that cheating, the new models left some questions unanswered. Who produces which differentiated product within a monopolistically competitive industry? The models, by construction, could not answer that question. We all invoked some notion of randomness, but without any explicit random mechanism in mind. Instead, what was crucially involved was a redefinition of the question. The detailed pattern of trade, the new theorists in effect argued, does not matter as long as aggregate measures like the volume of trade and the welfare effects of trade can be derived from the model. In effect, one had to step back from the blackboard and unfocus one’s eyes a bit, so as to grasp the broad pattern rather than the irrelevant details.

And once the new way of doing trade theory had been established, it made its way into pure comparative advantage modeling as well. Most notably, the important work of Eaton and Kortum (2002) on world trade patterns is based on a Ricardian model of comparative advantage – but the way it makes predictions about the parameters of a gravity equation that predicts the volume of bilateral trade in a multilateral world, rather than about specific goods imported and exported, is very much “new trade theory” in spirit.

Those who weren’t there, or haven’t participated in a comparable paradigm shift elsewhere in economics, probably can’t see why these “meta” changes in the way we did international trade theory were so hard, or made such a big difference. But they were, and they did. John Maynard Keynes famously described the process of arriving at his macroeconomic theory as a “journey of escape”; the emergence of the new trade theory was a similar journey, if less momentous, and had a profoundly liberating effect on the field.

4. GEOGRAPHY IGNORED

Bertil Ohlin's classic 1933 work laying out the beginnings of factor proportions theory was, of course, titled *Interregional and International Trade*. It has always been obvious that the motives for shipment of goods within countries are similar to those for shipment of goods between countries. It was also obvious, if one thought about it, that specialization within countries offered a new and possibly superior source of empirical evidence, if only because intra-national data are more likely to allow comparisons. So one might have expected the theory of international trade and the theory of economic geography to have developed in tandem, and in close relationship to each other, with a joint empirical research program.

In fact, however, as late as 1990 international economists took virtually no notice of trade within countries, or of the location of production in space. Nor was there really a strong independent presence of economic geography within the economics profession as a whole. Alfred Marshall (1890) may have devoted a whole chapter to "The concentration of specialized industries in particular localities", but that subject was barely touched on in the standard economics curriculum. Instead, if economic geography made any appearance at all, it was mainly urban economics, with a very limited discussion of location theory, with neither sub-field having any significant interaction with the much more well-established discussion of international trade.

Why was geography ignored by trade theorists? A large part of the explanation is the obvious centrality of increasing returns to geographical patterns: nobody really thinks that Silicon Valley owes its existence to exogenously given factors of production or Ricardian comparative advantage. (God made the Santa Clara valley for apricots, not semiconductors.) As long as trade theorists shied away from increasing returns in general, economic geography wasn't an inviting field.

Also, while there was elegant work in urban economics – I particularly admired Henderson's (1987) work on city systems – there was something deeply unsatisfying about the treatment of increasing returns in much of this literature. Essentially, the available techniques limited theorists to assuming external economies, leading to gibes that economists believed that agglomeration takes place as a result of agglomeration economies. A particular problem was that once one is simply assuming positive external economies, it's not at all clear how to think about the spatial limits of spillover. Do you have to be in the same city to reap positive externalities from other producers in the same industry? If so, why?

Over the course of the 1980s some researchers, notably Fujita (1988), realized that the monopolistic competition models could be used to derive endogenous externalities, explaining urban concentration. But these models depended on the assumption that the monopolistically competitive goods were completely nontradable, again offering little help on the question of the spatial reach of agglomeration economies.

But the analysis of the home market effect, already an established part of the new trade theory, suggested an approach to economic geography that did not depend on making goods strictly nontradable.

5. NEW ECONOMIC GEOGRAPHY

To lay out the logic of the “new economic geography” I find it helpful to talk in terms of a “cheat” model – a sort of model of the model – I originally devised to get some intuition about the home market effect. In this model we consider the location decision of a single producer serving two markets. We assume that the producer has fixed sales of S , S^* units of a good in the two markets, with $S > S^*$. And it must pay a transport cost of τ for each unit shipped from one location to the other. The producer has the option of having either one or two plants; by opening a second plant, the producer can eliminate transport costs but must pay an extra fixed cost F . Clearly, if the producer opens only one plant, it will be in the larger market. But will it concentrate production? Only if $F > \tau S^*$.

It goes without saying that this little exposition ignores market structure, pricing, the elasticity of demand, and more. But we know that we can put those things back in by doing a full Dixit-Stiglitz, and the cheat version conveys the essential intuition: if economies of scale, as captured by F/S^* , are large enough compared to transport costs, production will be geographically concentrated, and that concentration will, other things equal, be in the larger market.

From there it’s an obvious, short step (which for some reason took me a decade to take) to a model of geographical concentration of factors of production. Think now of a world in which there are many firms making the same kind of choice I just described, and also in which some but not all resources are mobile. Let S be the size of the overall market, μ be the share of that market attributable to “footloose” production, and suppose that there are two symmetric locations. Then we can think of a possible equilibrium in which all the footloose factors concentrate in one place. In that case the other location – the smaller market – would demand $S(1 - \mu)/2$ units of our representative good.

And this concentration of production would be self-sustaining if $F > \tau S(1 - \mu)/2$, or $F/S > \tau(1 - \mu)/2$. So that’s our criterion for the creation of a self-sustaining concentration of production in space.

As in the original home market effect exercise, this doesn’t quite get it right, because it fails to take account of market structure and demand elasticity. As I showed in Krugman (1991), it also misses a second reason for agglomeration: “forward linkages,” which in the simplest case would take the form of a lower cost of living for workers residing close to production concentrations. But the essential insights are right:

1. A self-sustaining concentration of production in space can occur if economies of scale (F/S) are large, transport costs low, and enough production is mobile.
2. Which location gets the concentration of production is arbitrary, and can be presumed to be a function of initial conditions or historical accident.

What was learned from this analysis? For one thing, it immediately cast light on some important aspects of economic history. Notably, the character of U.S. economic geography is well known to have gone through a sort of phase change in the middle of the 19th century, as the nation became differentiated into a farm belt and a manufacturing belt. What was happening at the time of this phase change? The rise of large scale production (economies of scale); railroads (lower transport costs); and a declining farm share in the economy (more mobile production). A simple model explains a major qualitative change in the economy.

This “core-periphery” model, essentially a model of agglomeration, was the starting point for the new economic geography. It was immediately clear, however, that one also wanted to model other key aspects of geography, notably regional specialization in different industries and the system of cities. With a few tricks, especially by assuming increasing returns in the production of intermediate as well as final goods, it turned out to be possible to address many of these issues.

As in the case of the new trade theory, a willingness to focus on tractable special cases was of the essence. In Fujita, Krugman, and Venables (1999) we described our method as “Dixit-Stiglitz, icebergs, evolution, and the computer”: we used Dixit-Stiglitz-type models to handle increasing returns and imperfect competition, “iceberg” transport costs that are proportional to fob prices, simple adaptive dynamics to arrive at equilibria, and numerical simulation to deal with models that tended to be just past the edge of paper-and-pencil analysis.

Also as in the case of the new trade theory, the new economic geography created a style of work that reached well beyond the specifics of the initial models. New economic geographers rediscovered Marshall’s chapter, which contained a beautifully laid out trinity of reasons for industry localization: knowledge spillovers (“the mysteries of the trade become no mysteries; but are, as it were, in the air”), labor market pooling, and specialized suppliers. Only the last of these three could be modeled with the original, new trade theory-inspired models, but the new focus on location led to reinvigorated interest in the others.

One gratifying result of the emergence of the new economic geography was a surge in empirical work. Some of this was driven specifically by the new models, but there was also a broader effect: the new models sensitized economists to the fact that regional variations in industrial specialization were an important laboratory for economic ideas, and that there was a great deal of evidence out there to be exploited. Before 1990 there were few high-profile economists using, say, cross-city comparisons to shed light on such subjects as

externalities, innovation, and growth. But after 1990 such studies exploded, and I believe that this was largely due to the emergence of the new economic geography.

With the emergence first of the new economic geography, then of a vast literature on external economies, the increasing returns revolution in trade and geography reached maturity. The old vision of the world economy, limited by the assumption of constant returns, had been superseded by an enlarged vision that incorporated the great tradition of trade theory but went well beyond it.

However, a funny thing happened on the way to this new vision. It's at least arguable that even as classical trade theory was being rejected or at least revamped, the world itself was becoming more classical, less driven by the increasing returns the new theory emphasized.

6. IS THE WORLD BECOMING MORE CLASSICAL?

Sometimes the progression of economic ideas mirrors changes in the real economy. Thus, macroeconomics emerged as a discipline at least in part because the business cycle became more severe over the first several decades of the 20th century. And the empirical observations that motivated the new trade theory largely concerned the rising role of increasing returns, as opposed to comparative advantage, in the growth of trade after 1950.

But there's no reason the world has to keep moving in a direction that makes new theories more relevant. And there's good reason to believe that the world economy has, over time, actually become less characterized by the kinds of increasing-returns effects emphasized by new trade and new geography.

In the case of geography, in fact, the peak impact of increasing returns probably occurred long before the new theorists arrived on the scene. Even in Krugman (1991a), I noted that the 1900 Census contained an extensive monograph on the localization of industries, emphasizing precisely the cumulative causation and role of historical accident that became central to new geography. And the history of such classic localizations as that of the car industry seemed, if anything, to suggest that concentrations due to increasing returns peaked before World War II. Meanwhile, the manufacturing belt itself began to dissolve after the war, and especially after the 1950s, as industry spread to the sunbelt.

Work by Kim (1998) seems to confirm the notion that the peak importance of increasing returns in industry location occurred circa 1930. Figure 4 shows his calculation of an index of regional manufacturing specialization (first proposed in Krugman, 1991a), using manufacturing censuses. From a peak in the interwar years, this index has since declined dramatically.




Figure 4. Regional specialization in manufacturing.

To be fair, this result may in part reflect statistical noise, as old industrial classifications fail to keep up with the modern division of labor. But the data accord with common perception: many of the traditional localizations of industry have declined (think of the Akron rubber industry), and those that have arisen, such as Silicon Valley, don't seem comparable in scale.

What about international trade? The rise of the new trade theory was motivated to a large extent by the rising relative importance of similar-similar trade: two-way exchanges of goods among advanced economies. For the last two decades, however, the trend has been in the other direction, with rapidly rising trade between advanced economies and much poorer, lower-wage economies, especially China. One simple indicator of this shift is the average hourly compensation of workers in top U.S. trading partners, as a percentage of U.S. compensation (Krugman, 2008). In 1975 that indicator stood at 76, and by 1990 it had risen slightly to 81, indicating that the United States was to a large extent trading with countries at a similar level of economic development. By 2005, however, the indicator had fallen to 65, largely reflecting the rapid growth in trade with China and Mexico. In 2006, for the first time, the United States did more trade in manufactured goods with developing countries than with other advanced nations.

And nobody doubts that trade between the United States and Mexico, where wages are only 13 percent of the U.S. level, or China, where they are only about 4 percent, reflects comparative advantage rather than arbitrary, scale-based specialization. The old trade theory has regained relevance.

Both new geography and new trade, then, may describe forces that are waning rather than gathering strength. Yet they're hardly irrelevant. And even the fact that they may be losing force is itself an important insight. For example, the contrast between the deep troubles of the Big Three automakers in the United States and the less afflicted foreign-owned operations, many of them located outside the traditional manufacturing belt, may in part reflect the diminishing advantages of being co-located with other producers in your industry.

Whether the influence of increasing returns on trade and geography is rising or falling, one thing is clear: much was learned from the intellectual revolution that brought increasing returns into the heart of how we think about

the world economy. It wasn't just that economists could make sense of previously puzzling data, we found ourselves able to see things that had previously been in an intellectual blind spot. Many people contributed to this process of enlightenment; I'm proud to have been a part of the journey.

REFERENCES

- Balassa, B. (1966), "Tariff Reductions and Trade in Manufactures among the Industrial Countries," *The American Economic Review*, Vol. 56, No. 3 (Jun., 1966), 466–473.
- Baldwin, R. and Martin, P. (1999), "Two waves of globalization: Superficial similarities, fundamental differences," National Bureau of Economic Research, Working Paper #6904.
- Burenstam Linder, S. (1961), *An essay on trade and transformation*. New York, John Wiley and Sons.
- Dixit, A. K. and Norman, V. (1980), *Theory of International Trade: A Dual, General Equilibrium Approach*, London: Cambridge University Press.
- Dixit, A. K. and Stiglitz, J. E. (1977), "Monopolistic competition and optimum product diversity," *American Economic Review*, 67:3, 297–308.
- Eaton, J. and Kortum, S. (2002), "Technology, Geography, and Trade," *Econometrica*, 70, 1741–1780.
- Fujita, M. (1988), "A monopolistic competition model of spatial agglomeration: A differentiated product approach," *Region. Sci. Urban Econ.* 18, 87–124.
- Fujita, M., Krugman, P., and Venables, A. (1999), *The spatial economy*, MIT Press.
- Grubel, H.G. and Lloyd, P.J. (1975), *Intra-Industry Trade: The Theory and Measurement of International Trade in Differentiated Products*, London: Macmillan.
- Helpman, E. (1981), "International Trade in the Presence of Product Differentiation, Economies of Scale, and Monopolistic Competition: A Chamberlin-Heckscher-Ohlin Approach," *Journal of International Economics*, 11, 305–340.
- Helpman, E. and Krugman, P. (1985), *Market Structure and Foreign Trade*, MIT Press.
- Henderson, J. V. (1987), "General equilibrium modelling of systems of cities," in Edwin S. Mills (ed.), *Handbook of Regional and Urban Economics*, 2 ed. Amsterdam: Elsevier, 927–956.
- Johnson, H. G. (1967), "International trade theory and monopolistic competition theory," in Kuenne, R. E. (ed.), *Monopolistic Competition Theory: Studies in Impact; Essays in Honor of Edward H. Chamberlin*, New York: John Wiley, 203–218.
- Krugman, P. R. (1979), "Increasing returns, monopolistic competition and international trade," *Journal of International Economics*, 9, 469–479.
- Krugman, P. R. (1980), "Scale economies, product differentiation and the pattern of trade," *American Economic Review*, 70, 950–959.
- Krugman, P. R. (1981), "Intra-industry specialization and the gains from trade," *Journal of Political Economy*, 89, 959–973.
- Krugman, P. R. (1989), "Industrial organization and international trade," in R. Schmalensee and R. Willig (eds.), *Handbook of Industrial Organisation*, Vol. 2, Amsterdam: North-Holland, 1179–1223.
- Krugman, P. R. (1991a), *Geography and Trade*, Cambridge: MIT Press.
- Krugman, P. R. (1991b), "Increasing returns and economic geography," *Journal of Political Economy*, 99, 483–499.
- Krugman, P. (2008), "Trade and wages, reconsidered," *Brookings Papers on Economic Activity*.
- Lancaster, K. J. (1979), *Variety, Equity and Efficiency*, New York: Columbia University Press.
- Lancaster, K. J. (1980), "Intra-industry trade under perfect monopolistic competition," *Journal of International Economics*, 10, 151–175.
- Marshall, A. (1890), *Principles of Economics*, London: Macmillan.
- Meyer, D. (1983), "Emergence of the American Manufacturing Belt: An Interpretation," *Journal of Historical Geography* 9, 145–174.
- Norman, V. (1976), "Product differentiation and trade," ms., U.K. Economic Theory Study Group, University of Warwick.
- Spence, A. M. (1976), "Product selection, fixed costs and monopolistic competition," *Review of Economic Studies*, 43, 217–236.
- Portrait photo of Paul Krugman by photographer Ulla Montan.