

# Space: The Final Frontier

Paul Krugman

**T**he title of this article was suggested by the back of a T-shirt, which I received as a gift from students at the University of Pennsylvania's regrettably vanished Department of Regional Science. On the shirt's front was a portrait, not of James T. Kirk, but of Walter Isard, who for many years—most notably in his 1956 book *Location and Space-Economy*—tried to get his fellow economists to give the spatial aspects of the economy their proper due. Despite Isard's efforts, however, one must say that until the early 1990s the traditional neglect of spatial economics, of where economic activity takes place and why, remained pretty much intact. Even now not one of the best-selling introductory textbooks in economics contains a single index entry for "location," "space," or "regions." (Most do not even contain an entry for "cities").

In the last six or seven years, however, interest in spatial economics has surged. In this article I will try to summarize briefly the reasons for that surge; the key elements of the so-called "new economic geography;" the current state of research; and the prospects and difficulties facing this subfield of economics.

## The Neglect of Space

In a way, there is no puzzle about why economists have turned their attention to the location of economic activity: it is inherently a very interesting subject. (Ask any real estate developer.) The question is why it took them so long,

■ *Paul Krugman is Professor of Economics, Massachusetts Institute of Technology, Cambridge Massachusetts. His e-mail address is <krugman@mit.edu>.*

and this question can perhaps best be answered by describing a particular, important example of a good idea that was neglected by all but a handful of economists for decades.

In the 1950s and 1960s geographers—including especially Harris (1954) and Pred (1966)—developed what in retrospect can be seen as a coherent and intuitively compelling story about the emergence of large regional concentrations of economic activity, such as the “manufacturing belt” in America’s northeast and inner midwest.

Harris’s (1954) idea was a straightforward one: producers will tend, other things being the same, to choose sites with good access to markets. He measured the market access of each county in the United States using several ad hoc but fairly sensible measures of “market potential”: a weighted sum of purchasing power across locations, with the weight for each location depending inversely on its distance. His results showed that the heavily industrialized regions of the United States were in general also locations with exceptionally high market potential. This was not too surprising: precisely because a large part of U.S. population and production was concentrated in the manufacturing belt, locations in that belt had better market access than locations elsewhere in the country. But this simple observation led Harris to an interesting, even exciting suggestion: that the concentration of production was self-reinforcing. Firms chose to produce in regions with good access to markets; but access to markets tended to be good in regions in which many firms chose to produce.

Where Harris (1954) looked at a snapshot of the economy at a point in time, Pred (1966) was interested in the dynamics of regional growth. It was then and still is common for regional planners to work with a simple “base-multiplier” model of regional income. That is, they start with a projection of the “export” earnings of a region (its sales to other regions inside and outside of the country), then use an estimate of the share of income spent within the region to compute a multiplier on that base. Thus if one projects that export income will be \$10 billion, and that 60 percent of income will be spent locally, one would project regional income of  $10/(1-.6) = \$25$  billion. Pred argued, however, that both the size of the export base and the share of income spent locally were likely to be increasing functions of the size of the regional economy—and that this meant that if a regional economy for whatever reason reached a sufficiently large scale, it could take off in a cumulative process of growth. For example, the large market might make it profitable to produce locally goods that had previously been imported from other regions; this would increase the multiplier on the region’s export base, leading to a further expansion of income, which would lead to still more local production, and so on. Although it is unclear to what extent Pred and others realized this (I have made these ideas sound much clearer than they are either in the original expositions or in the traditional regional science literature; textbooks like Dicken and Lloyd (1990) treat Harris and Pred in different chapters, and do not highlight the similarity of what they had to say), it is very close to Harris’s story: again, a large market

attracts producers, which further increases the size of that market, which attracts still more producers, and so on.

While one might raise empirical questions about the strength of the posited linkages, this is an intuitively plausible explanation of something of which we are all aware: the dramatically uneven distribution of economic activity across space. Furthermore, it might seem on the face of it to be exactly the kind of story an economist would love; that is, it is an account of how the actions of self-interested individual agents interact to produce aggregate behavior that is more than the sum of its parts. And it is an idea that is simple enough to be explained to almost anyone: college freshmen, reporters, even politicians. Yet economists did not pick up on the idea until around 1990.<sup>1</sup> Why?

The most likely answer is that underlying the work of Harris (1954) and Pred (1966) is the implicit assumption that there are substantial economies of scale at the level of the plant. In the absence of such scale economies, producers would have no incentive to concentrate their activity at all: they would simply supply consumers from many local plants, vitiating Harris's argument. And expansion of a regional market would not predictably lead to any increase in the range of goods produced within that region, vitiating Pred's. Increasing returns, in other words, were central to the story.

The same may be said of spatial economics in general. Almost all of the interesting ideas in location theory rely implicitly or explicitly on the assumption that there are important economies of scale enforcing the geographic concentration of some activities. Thus Weber's (1909) analysis of the location decisions of an individual producer trying to minimize the combined costs of producing and delivering his product assumes that there can be only one production site; Christaller's (1933) suggestion that cities form a hierarchy of central places depends on the assumption that larger cities can support a wider range of activities; Lösch's (1940) famous demonstration that an efficient pattern of central places would imply hexagonal market areas assumes that there are economic activities that can be undertaken only at a limited number of sites. (The main example of a location model that does not rely on some form of scale economies, the land-rent analysis of von Thunen (1826), in effect hides the role of increasing returns by simply assuming the existence of a central city.) But unexhausted economies of scale at the level of the firm necessarily undermine perfect competition; and in the 1950s and 1960s there were no tractable models of imperfect competition. (Lösch and Christaller both seem to be describing planning solutions rather than market outcomes.) As a result, the

<sup>1</sup> Such an assertion is always risky. While I am not aware of any pre-1990 articles by economists that clearly discuss the circular logic suggested by Harris and Pred, I have not made a serious scholarly search of the literature. But certainly the idea did not make it into mainstream economic discussion, even among economists who did work on locational issues; for example, there is no sign of that logic in the papers collected in Karaska and Bramhall (1969).

mainstream economic theory of the time turned a blind eye to most stories, no matter how simple or intuitive, for which increasing returns were crucial.

The reason why space has finally made it into the economic mainstream is therefore obvious: imperfect competition is no longer regarded as impossible to model, and so stories that crucially involve unexhausted scale economies are no longer out of bounds. Indeed, the new interest in space may be regarded as the fourth (and final?) wave of the increasing returns/imperfect competition revolution that has swept through economics over the past two decades. First came the New Industrial Organization, which created a toolbox of tractable if not entirely convincing models of imperfect competition; then the New Trade Theory, which used that toolbox to build models of international trade in the presence of increasing returns; then the New Growth Theory, which did much the same for economic growth. What happened after 1990 was the emergence of the New Economic Geography, which might perhaps be best described as a “genre”: a style of economic analysis which tries to explain the spatial structure of the economy using certain technical tricks to produce models in which there are increasing returns and markets are characterized by imperfect competition.

### **The New Economic Geography: Tricks of the Genre**

Because the main obstacle that economists have traditionally faced when trying to confront issues involving increasing returns is one of tractability, overcoming that obstacle depends crucially on technical tricks: on strategic assumptions that may be unrealistic but make a model easier to build, on clever new ways of solving models that might otherwise seem too complex to deal with. To date, the new economic geography has depended heavily on the tricks summarized in Fujita, Krugman and Venables (forthcoming) with the slogan “Dixit-Stiglitz, icebergs, evolution, and the computer.” Let us consider each part of that slogan in turn.

*Dixit-Stiglitz:* The remarkable model of monopolistic competition developed by Dixit and Stiglitz (1977) has become a workhorse in many areas of economics. In the new economic geography, it has one especially appealing feature: because it assumes a continuum of goods, it lets the modeler respect the integer nature of many location decisions—no fractional plants allowed—yet analyze the model in terms of the behavior of continuous variables like the share of manufacturing in a particular region. In effect, Dixit-Stiglitz lets us have our cake and cut it into arbitrarily small pieces, too. The price of that convenience is, of course, that Dixit-Stiglitz is a very restrictive, indeed in some respects, silly model. Above all, the assumed symmetry among varieties, and the resulting absence both of monopoly rents in equilibrium and of any strategic behavior by firms, mean that Dixit-Stiglitz analyses undoubtedly miss much of what really happens in imperfectly competitive industries.

*Icebergs:* This is a less familiar technical trick. In location theory, transportation costs are of the essence; yet any attempt to develop a general-equilibrium model of

economic geography would be substantially complicated by the need to model the transportation as well as the goods-producing sectors. Worse yet, transportation costs can undermine the constant demand elasticity that is one of the crucial simplifying assumptions of the Dixit-Stiglitz model. Both problems can be sidestepped with an assumption first introduced by Paul Samuelson (1952) in international trade theory: that a fraction of any good shipped simply “melts away” in transit, so that transport costs are in effect incurred in the good shipped. In the new geography models, melting is usually assumed to take place at a constant rate per distance covered; for example, 1 percent of the cargo melts away per mile. In terms of modeling convenience, there turns out to be a spectacular synergy between Dixit-Stiglitz market structure and “iceberg” transport costs: not only can one avoid the need to model an additional industry, but because the transport cost between any two locations is always a constant fraction of the f.o.b. price, the constant elasticity of demand is preserved. It is too bad that actual transport costs look nothing like that.

*Evolution:* Interesting stories about economic geography often seem to imply multiple equilibria. Suppose, for example, that you follow Harris and Pred in arguing that producers want to locate where other producers choose to locate; this immediately suggests some arbitrariness about where they actually end up. But which equilibrium does the economy select? New economic geography models typically assume an ad hoc process of adjustment in which factors of production move gradually toward locations that offer higher current real returns. This sort of dynamic process was initially proposed apologetically, since it neglects the role of expectations. But it is possible to regard models of geography as games in which actors choose locations rather than strategies—or rather in which locations *are* strategies—in which case one is engaged not in old-fashioned static expectations analysis but rather in state-of-the-art evolutionary game theory! (To middle-brow modelers like myself, it sometimes seems that the main contribution of evolutionary game theory has been to re-legitimize those little arrows we always wanted to draw on our diagrams.) Indeed, the ad hoc adjustment rules introduced in Krugman (1993a) are, though I did not know it at the time, precisely the rules of “replicator dynamics” (see, for example, Weibull, 1995).

*The computer:* Finally, despite the best efforts of the theorist, all but the simplest models of economic geography usually turn out to be a bit beyond the reach of paper-and-pencil analysis. As a result, the genre relies to an unusual extent on numerical examples—on the exploration of models using both static calculations and dynamic simulations.

All of these elements of the new economic geography, including its weaknesses, were present in the genre’s founding paper, Krugman (1991a). That paper was, in effect, an attempt to formalize the story suggested by Harris (1954) and Pred (1966). The model envisaged an economy consisting of two regions (although the setup allowed easy extension to multi-regional settings), with two industries: perfectly competitive agriculture and imperfectly competitive (Dixit-Stiglitz) manufac-

turing. The agricultural good was produced by immobile farmers;<sup>2</sup> but manufactures were produced by workers who could move to the region that offered the higher real wage. And manufactures (but not agriculture) were subject to iceberg transportation cost.

Although the economies of scale assumed in the model were purely internal to firms, “centripetal” forces—forces that tend to make manufacturing concentrate in only one region, and act a lot like external economies—emerge from the three-way interaction among scale economies, transportation costs, and factor mobility. Loosely speaking, firms want to *concentrate* production (because of scale economies) *near* markets and suppliers (because of transport costs); but access to markets and suppliers is best where *other firms locate* (because of market-size effects). This circular logic can produce agglomerations—although it is opposed by the “centrifugal” force generated by the immobility of agriculture, which provides an offsetting incentive to locate in the region with fewer local competitors.

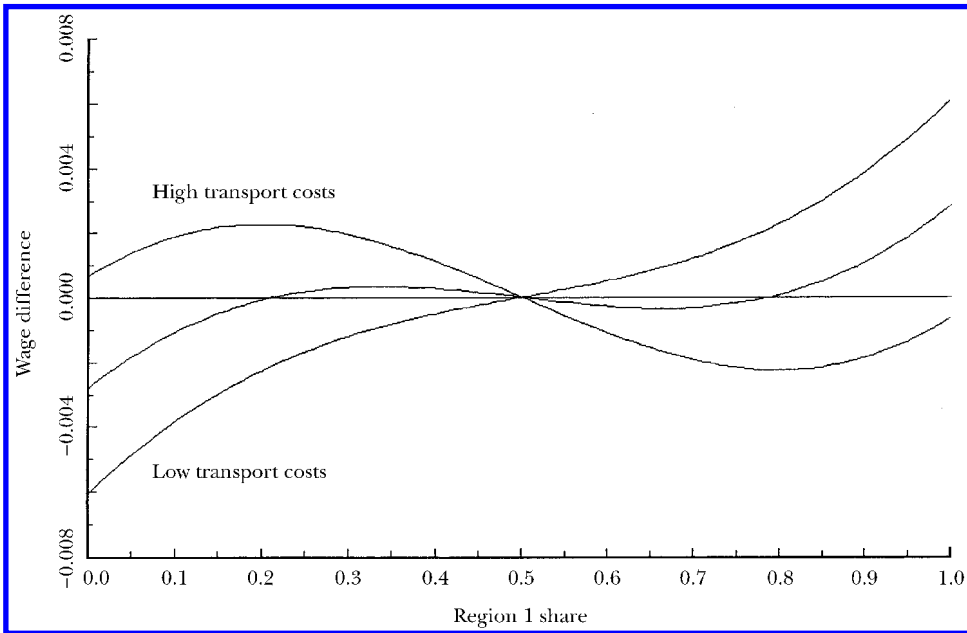
This most basic model can, with some effort, be largely solved analytically; but it is easiest to examine via numerical examples. In particular, it turned out to be illuminating to calculate how the difference in real wages between regions depending on the allocation of manufacturing between those regions (a calculation that involves repeatedly solving a small computable general equilibrium model). Figure 1 shows a typical set of calculations. On the horizontal axis is the share of the population of workers in region 1; on the vertical axis the difference between the real wage in region 1 and that in region 2. Each curve is calculated for a different level of transport costs.

The rough intuition behind these curves runs as follows. In the case of high transport costs, there is relatively little interregional trade; so the wages workers can earn depend mainly on the amount of local competition, and are thus decreasing in the number of other workers in the same region. On the other hand, when transport costs are low, a typical firm sells extensively in both regions; but since it has better access to markets if it is located in the region with the larger population, it can afford to pay higher wages—and the purchasing power of these wages is also higher because workers have better access to consumer goods. So in that case real wages are increasing in a region’s population. As a result, concentration of population in either region is an equilibrium, since no individual worker will have an incentive to move. At intermediate transport costs these two forces are nearly balanced. The particular curvature shown, in which centripetal forces are stronger when regions are very unequal, while centrifugal forces are stronger when they are nearly symmetric, is an artifact of the particular functional forms used in this exercise; Fujita, Krugman and Venables (forthcoming) shows that the curvature can be reversed in closely related models.

<sup>2</sup> Why assume immobile farmers rather than have land as a separate factor and allow labor to move between agriculture and manufacturing? Essentially for technical convenience: adding land-labor substitution substantially complicates the model without adding much in the way of insight.

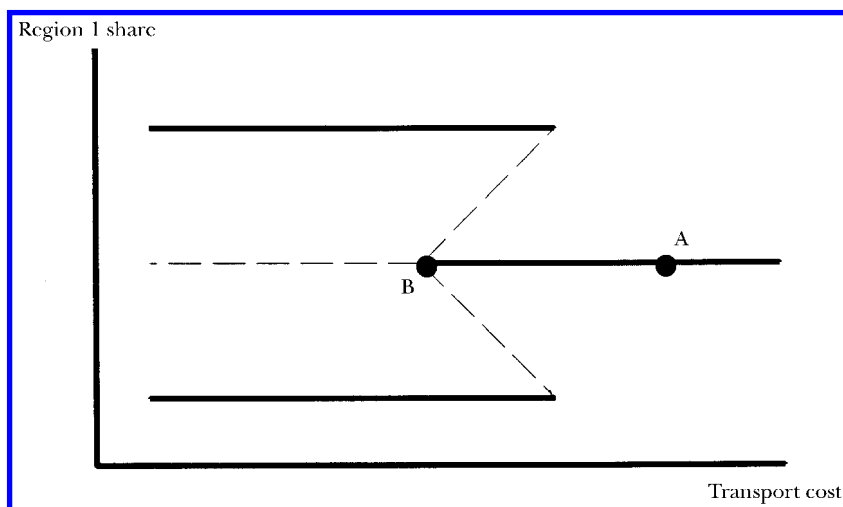
Figure 1

## Population Shares and Wage Differences Between Two Regions



Since workers are assumed to move to whichever region offers the higher real wage, in the case of high transport costs there is a unique equilibrium with workers evenly divided between the regions. In the case of low transport costs there are three equilibria: one with workers evenly divided, one with workers concentrated in either region. And in the intermediate case there are five equilibria.

The figure immediately suggests the appeal of some sort of evolutionary dynamics—which in this case simply amount to the assumption that workers move gradually toward the region offering the higher real wage—as a way of declaring some of these equilibria unstable. If one allows that assumption, it is possible to use simple bifurcation diagrams to summarize how the economy's behavior depends on its parameters. For example, Figure 2 shows how the set of equilibria (as measured by the share of the manufacturing labor force in region 1) depends on transport costs, with solid lines indicating stable and broken lines unstable equilibria. The figure illustrates nicely one of the appealing features of the new economic geography: it easily allows one to work through interesting “imaginary histories.” Suppose, for example, that we imagine an economy that starts with high transport costs and therefore with an even division of manufacturing between regions, a situation illustrated by the point labeled *A*. Then suppose that transport costs were gradually to fall. When the economy reached *B*, it would begin a Pred-type process,

*Figure 2***Bifurcation into Core and Periphery**

in which a growing concentration of manufacturing in one region would lead to a still larger concentration of manufacturing in that region. Which region ends up as the manufacturing “core,” which as the agricultural “periphery,” would presumably depend on small historical accidents. That is, the economy would spontaneously organize itself into a core-periphery geometry strongly suggestive of the actual division of the United States into manufacturing and farm belts—or perhaps of the division of Italy into prosperous north and impoverished south.

In short, the technical tricks that characterize the new economic geography have opened the door to theoretical modeling of the sort of inherently exciting spatial economic issues that economists had previously neglected. Let me turn next to a survey of some of the issues addressed by this rapidly growing literature.

### **Regions, Cities, and Nations**

Broadly speaking, theoretical work in the new economic geography has moved in two directions. One direction has been an effort to build links from the new genre to traditional questions of location theory. The other has been an effort to use the genre as the basis for a new, “spatial” view of international trade.

If you want to use the new economic geography to bring the grand tradition of location theory into the economic mainstream, you are likely to be unsatisfied with a two-region model, which does not have much spatial content. What you really want is a multi-location model, or even better, a model with continuous space. As



long as you are willing to rely on numerical examples, however, the new economic geography style of model can easily handle any number of regions, with whatever “geometry” of transport costs one likes.

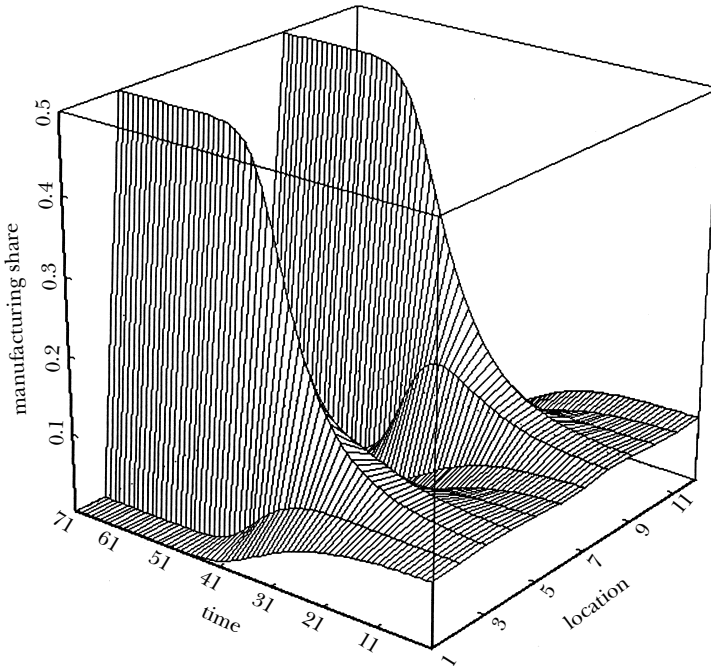
An interesting aspect of such multi-regional models is that they provide a justification for a version of the “market potential” function used by Harris (1954). In this case, market potential can be defined as the real wage manufacturing firms can afford to pay in any given location; it is a function of access to markets, although not quite as simple a function as Harris used. And the dynamics of the economy can, if one likes, be viewed as the coevolution of two landscapes: a landscape of current distribution of economic activity, which determines a second landscape of market potential, which in turn determines how that first landscape changes over time.

The results of multi-region simulations depend, of course, not only on parameters but on the geometry of the economy. An interesting if artificial special case has turned out to be the “racetrack” economy: an economy whose regions are laid out around a circle, with transportation possible only along that circle’s circumference. In such an economy a uniform distribution of manufacturing is always an equilibrium, referred to in Fujita, Krugman and Venables (forthcoming) as the Flat Earth. However, the Flat Earth may be unstable: the circular logic of concentration can cause an even slightly perturbed Flat Earth to spontaneously develop one or more local concentrations of manufacturing. Figure 3 (which Robert Gordon instantly dubbed the “’59 Cadillac diagram” at a seminar presentation) illustrates such a process for a 12-region racetrack model (in which region 12 is next to region 1): from an almost flat but randomly perturbed initial distribution of manufacturing, the model evolves into a structure in which all manufacturing is concentrated in only two regions. In this example the two “winning” regions are exactly opposite one another—an unexpected regularity that consistently emerges in such examples as long as the initial distribution of manufacturing is sufficiently flat.<sup>3</sup>

An alternative way to go beyond two-region modeling is to use the new economic geography to revisit some of the questions of traditional location theory. In a series of papers, Mashisa Fujita and his students have in essence tried to take the German tradition of urban modeling that began with von Thünen (1826) and give it a true microeconomic foundation. (In Fujita’s models all labor is mobile; thus

<sup>3</sup> The reasons for this regularity, it turns out, can be understood using an approach originally suggested by none other than Alan Turing (1952). Turing was interested in the interacting effects of chemical signals diffusing around a ring of cells, but his approach works equally well for the distribution of manufacturing around a ring of regions. It involves linearizing the model in the vicinity of the Flat Earth, then representing the initial distribution of manufacturing (or whatever) as a Fourier series. It turns out that the components of that Fourier representation are also eigenvectors of the linearized model, so that you can, in effect, think of the components of that series as growing independently—and of the distribution becoming increasingly dominated by a fluctuation at some “preferred frequency” that depends on the parameters of the model, but *not* on the initial conditions. And it is this preferred frequency that determines the eventual number of agglomerations that emerge. (This is only one example of the tantalizing affinity that one often finds between the new economic geography and fashionable scientific trends like “complexity” theory.)

Figure 3

**Perturbing a 12-Region Racetrack Model**

the location of agriculture as well as manufacturing is endogenous.) In Fujita and Krugman (1995), a version of the original von Thünen model is offered in which the existence of a central city is no longer simply assumed; instead, manufacturing concentrates in the city because of the forward and backward linkages generated by that very concentration. Or looking at the issue another way, the concentration of economic mass at the city generates a self-validating cusp at that point in the market potential function that determines where manufacturing locates. Agriculture is then spread around that center, with land rents declining to zero at the agricultural frontier. Such a monocentric equilibrium, however, turns out to be sustainable only if the population is sufficiently small. Fujita and Mori (1996a) take the same basic model, but envision a gradually rising population which leads to the periodic emergence of new cities in a “long, narrow” economy that gradually spreads along a line; the resulting multi-city spatial economy may be regarded as a (one-dimensional) version of Lösch’s central place theory. (Nobody has yet managed to produce a model with Lösch’s famous hexagonal market areas.) Fujita, Krugman, and Mori (1997) consider an economy with multiple manufacturing industries, differing in transport costs and/or scale economies; such an economy spontaneously develops a system of central places that finally provides a justification

(again in only one dimension) for Christaller's (1933) hierarchical model of central places. Finally, Fujita and Mori (1996b) address an issue that, astonishingly, appears never to have been formally modeled before: the reasons why ports and other transportation nodes so often become the sites of major cities. What they show is that such transport nodes generate cusps in the market potential function, and therefore tend to serve as the seeds for city growth.

Moving from the local to the global, Anthony Venables and his students have tried to use new economic geography models as the basis for a new style of international trade model. What is the difference between regions and countries? One answer is that factors of production are far less mobile between countries than between regions of the same country, and in Venables-type models they are normally assumed to be completely immobile. Nonetheless, Venables (1995) shows that a circular process leading to economic differentiation between nations can still result if there are intermediate goods produced with economies of scale and subject to transport costs. In that case, a country with a large manufacturing sector offers a large market for intermediates; this leads to a concentration of intermediate production in that country, which gives it a cost advantage in downstream production, which further reinforces its advantage, and so on. In Krugman and Venables (1995) this story is used as the basis for a "history of the world," in which gradually declining transport costs lead first to the spontaneous differentiation of the world into a high-wage industrial "core" and a low-wage agricultural "periphery," then to a later convergence of wages as the periphery industrializes. Puga and Venables (1996) offer an alternative version of this story in which the driving force is the growing size of the market rather than growing economic integration. And Krugman and Venables (1997) use the "racetrack" geometry to model global international trade and specialization in a world in which borders are irrelevant, and in which even economic regions are left unspecified; nonetheless, the world spontaneously organizes itself into manufacturing zones surrounded by agricultural hinterlands.

Most recently, Baldwin and Forskild (1997) have developed an alternative version of geography and trade analysis, this time relying on tools borrowed from the endogenous growth literature. In these models the circular process involves not the movement of factors but their accumulation: countries with large markets invest more, which further enlarges those markets.

It is also possible to mix trade and urban economics using the new genre. Krugman and Livas (1996), for example, develop a model suggested by the relative decline of Mexico City as Mexico has opened itself to trade; the basic idea is that the importance of access to domestic consumers and suppliers, which was crucial as long as Mexico pursued a policy of inward-looking industrialization, has become much less relevant now that it exports more of its output and imports more of its intermediate goods.

As this partial survey of the theory indicates, the new economic geography has opened the door for analytical discussion of interesting and important issues that were, for the most part, previously ignored by economists. That in itself is something

of an achievement. Moreover, partly because of the novelty, partly because of the inherent sexiness of the stories—geography models naturally produce multiple equilibria, dramatically “catastrophic” changes in outcomes from small changes in parameters, large effects from small differences in initial conditions, and spontaneous emergence of unexpected order from randomness—these models are a lot of fun to work with.

But are they really relevant?

## **The Empirical Frontier**

An unfortunate feature of much of the “new” theorizing since the 1970s is that it has failed to lead to much validating empirical work. The new industrial organization has been notoriously better at creating interesting models than at generating empirical predictions; the new growth theory gave rise to a massive industry of cross-country growth regressions, but with few exceptions, these regressions have neither been closely tied to the theory nor provided clear evidence in its support. Under the combined influence of the new growth and new geography movements, there has been a parallel effort to extract insight from cross-sectional regressions on the growth of metropolitan areas; with only a few exceptions like Ades and Glaser (1997), however, these studies have similarly failed to offer much direct testing of the specifics of the models.

Perhaps the closest thing to a direct test of the models has been the recent work of Davis and Weinstein (1996, 1997), who have used international and inter-regional production and consumption data to test for the “home market effect”: the prediction, made by the underlying models of new geography, that a larger demand for the products of an industry in any given region will lead, other things being the same, to a more than one-for-one increase in the regional production of that industry. Their results are generally negative for international comparisons, but generally positive at the regional level.

It would not be surprising if it turns out that the market-size effects emphasized by the current generation of new geography models are a less important source of agglomeration, at least at the level of urban areas, than other kinds of external economies. It is, for example, a well-documented empirical regularity that both plants and firms in large cities tend to be *smaller* than those in small cities (for example, Hoover and Vernon, 1959); this suggests that big cities may be sustained by increasing returns that are due to thick labor markets, or to localized knowledge spillovers, rather than those that emerge from the interaction of transport costs and scale economies at the plant level. However, serious empirical work—which will probably require detailed micro studies of particular industries—still remains to be carried out.

Eventually, one might hope to develop “computable geographical equilib-

rium'' models, which can be used to predict the effects of policy changes, technological shocks, etc. on the economy's spatial structure in the same way that such models are currently used to predict the effects of changes in taxes and trade policy on the economy's industrial structure. However, preliminary efforts in this direction by several researchers, myself included, have found that such models are not at all easy to calibrate to actual data; in general, the tendency toward agglomeration is stronger in the models than it seems to be in the real economy!<sup>4</sup>

At this point, then, the new economic geography—like its sister genre, the new growth theory—has been more successful at raising questions than at answering them, better at creating a language with which to discuss issues than at creating the tools to resolve those discussions. Nonetheless, it is better to have an unresolved discussion about the spatial aspects of the economy than to ignore them. The economy exists in and occupies space on the map; it is good that we have finally begun to acknowledge this in our models.

<sup>4</sup>One abortive effort of mine was an attempt to build a model of the distribution of industry in the United States circa 1910. It turned out that no matter how I fiddled with the parameters, the model offered only two alternatives. If I chose parameters that made the advantages of the established industrial concentrations strong, California refused to develop any industry; if I made them weak, California would industrialize, but so would Iowa.

## References

- Ades, A. and E. Glaser**, "Trade and circuses: explaining urban giants," *Quarterly Journal of Economics*, 1995, 110:1, 195–227.
- Christaller, W.**, *Central Places in Southern Germany*. Jena: Fischer, 1933.
- Davis, D. and R. Weinstein**, "Empirical testing of economic geography: evidence from regional data," mimeo: Harvard University, 1997.
- Dicken, P. and P. Lloyd**, *Location in Space: Theoretical Perspectives in Economic Geography*. New York: HarperCollins, 1990.
- Dixit, A. and J. Stiglitz**, "Monopolistic competition and optimum product diversity," *American Economic Review*, 1977, 67:3, 297–308.
- Fujita, M. and P. Krugman**, "When is the economy monocentric: von Thunen and Chamberlin unified," *Regional Science and Urban Economics*, 1995, 25:4, 505–28.
- Fujita, M., P. Krugman, and A. Venables**, *The Spatial Economy*, forthcoming.
- Fujita, M. and T. Mori**, "Structural stability and evolution of urban systems," *Regional Science and Urban Economics*, 1996a, 27:4–5, 399–442.
- Fujita, M. and T. Mori**, "The role of ports in the making of major cities: self-agglomeration and hub-effect," *Journal of Development Economics*, 1996b, 49:1, 93–120.
- Fujita, M., T. Mori, P. Krugman**, "On the evolution of hierarchical urban systems", mimeo, Kyoto, 1997.

- Harris, C.D.**, "The market as a factor in the localization of production," *Annals of the Association of American Geographers*, 1954, 44, 315–348.
- Hoover, E. and R. Vernon**, *Anatomy of a Metropolis*. Cambridge: Harvard, 1959.
- Karaska, G. and D. Bramhall, eds.**, *Locational Analysis for Manufacturing*. Cambridge: MIT Press, 1969.
- Krugman, P.**, "Increasing returns and economic geography," *Journal of Political Economy*, 1991a, 99:3, 483–99.
- Krugman, P.**, *Geography and Trade*. Cambridge: MIT Press, 1991b.
- Krugman, P.**, "On the number and location of cities," *European Economic Review*, 1993a, 37:2–3, 293–98.
- Krugman, P. and R. Livas Elizondo**, "Trade policy and the third world metropolis," *Journal of Development Economics*, 1996, 49:1, 137–50.
- Krugman, P. and A. Venables**, "Globalization and the inequality of nations," *Quarterly Journal of Economics*, 1995, 110, 857–80.
- Krugman, P. and A. Venables**, "The seamless world: a spatial model of international specialization and trade," mimeo, MIT, 1997.
- Lösch, A.**, *The Economics of Location*. Jena: Fischer, 1940.
- Pred, A.R.**, *The Spatial Dynamics of U.S. Urban-Industrial Growth, 1800–1914*. Cambridge: M.I.T. Press, 1966.
- Puga, D. and A. Venables**, "The spread of industry: spatial agglomeration in economic development," CEPR Working Paper #1354, 1997.
- Samuelson, P.**, "The transfer problem and transport costs," *Economic Journal*, 1952, 64, 264–89.
- von Thünen, J.**, *The Isolated State*. English edition, London: Pergamon, 1826.
- Turing, A.**, "The chemical basis of morphogenesis," *Philosophical Transactions of the Royal Society of London*, 1952, 237, 37.
- Venables, A.**, "Equilibrium locations of vertically linked industries," *International Economic Review*, 1996, 37:2, 341–59.
- Weber, A.**, *Theory of the Location of Industries*. Chicago: University of Chicago Press, 1909.
- Weibull, M.**, *Evolutionary Game Theory*. Cambridge: MIT Press, 1995.