

NBER WORKING PAPER SERIES

THE SEAMLESS WORLD: A SPATIAL
MODEL OF INTERNATIONAL
SPECIALIZATION

Paul Krugman
Anthony J. Venables

Working Paper No. 5220

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 1995

This paper is part of NBER's research programs in International Finance and Macroeconomics, and International Trade and Investment. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

© 1995 by Paul Krugman and Anthony J. Venables. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

THE SEAMLESS WORLD: A SPATIAL
MODEL OF INTERNATIONAL
SPECIALIZATION

ABSTRACT

This paper is an effort to do international trade theory without mentioning countries. Nearly all models of the international economy assume that trade takes place between nations or regions which are themselves dimensionless points. We develop a model in which economic space is instead assumed to be continuous, and in which this "seamless world" spontaneously organizes itself into industrial and agricultural zones because of the tension between forces of agglomeration and disagglomeration. One might expect such a model to be analytically intractable, but we are able to gain considerable insight through a combination of simulations and an analytical approach originally suggested in a biological context by Alan Turing.

Paul Krugman
Department of Economics
Stanford University
Stanford, CA 94305
and NBER

Anthony J. Venables
Department of Economics
London School of Economics
Houghton Street
London WC2A 2AE
UNITED KINGDOM

This paper is an attempt to do international trade theory without mentioning countries.

In recent years it has become common to argue that external economies operate at a local rather than a national level (Porter 1990), an observation that has led to a new interest in economic geography among international trade theorists (Krugman 1991). Taken to its logical conclusion, this line of thought seems to suggest an approach to international economics in which one thinks of countries as areas on a surface rather than as discrete locations, and tries to adopt a "field theory" approach to global trade that describes all flows of goods and services across that surface rather than only those flows that happen to cross arbitrary lines called borders. One might admit that in practice economic regions are often separated by natural barriers such as oceans and mountain ranges, and also that national borders appear to matter a great deal for trade flows, even when formal trade barriers are low. Nonetheless, it seems obvious that it would be useful to have at least a benchmark model of specialization and trade in a "seamless" world, one in which national boundaries are ignored and even economic regions -- which will typically be blobs with fuzzy edges rather than points -- are observed rather than defined ex ante.

We are not, however, aware of any formal models of world specialization and trade that do away with the assumption that regions, whether or not they are identified with nations, are discrete

points. Even in our own work on the emergence of global specialization (Venables (1993), Krugman and Venables forthcoming) we have focussed on two-region models in which trade between the regions is costly but trade within them is not -- that is, we have represented regions as exogenously defined points, even if the role these regions end up taking in the world economy is endogenous. There is, of course, no mystery about why international trade theorists have been reluctant to think in terms of a seamless world: at first sight, modeling such a world appears to be a task of daunting technical difficulty. After all, a three-location model is far more complex than a two-location model, a four-location model more complex still; surely, one supposes, a model with continuous space, which is to say an infinite number of locations, will be entirely intractable.

But then again, maybe not. In this paper we show that it is possible to gain some surprisingly clear intuitions about specialization and trade in a seamless world. Admittedly, these results depend on the use of a highly stylized dynamic model, in which we make extremely unrealistic assumptions about both natural geography and the motives for trade. Even within this stylized framework, we are obliged to rely to a considerable extent on simulation -- that is, on high-tech numerical examples. We are, however, able to develop intuition about these simulations and even derive a number of analytical results using insights from an

unexpected source: Alan Turing's classic analysis of the dynamics of morphogenesis (Turing 1952). On reflection, the idea that the origins of spatial structure in biology and economics might be analyzed using similar tools may not seem that outlandish; in any case, one purpose of this paper is to introduce "Turing waves", long a familiar concept in theoretical biology, to the economics literature. (See also the discussion in Krugman (forthcoming)).

This paper, then, is in five parts. Part 1 lays out the basic concepts, and offers an intuitive overview of the paper's analytical approach. Part 2 lays out the model, and shows how short-run equilibrium is determined. Part 3 shows, through simulation and then an analytical interpretation of those simulations, how a world with an initial distribution of economic activity that is (almost) uniform across space can spontaneously differentiate itself into industrial and agricultural regions. Part 4 describes how growing integration of such a structured world economy leads to "bifurcations", points at which the existing geographical structure becomes unstable and a new structure evolves. Finally, Part 5 summarizes the results and offers suggestions for further research.

1. An intuitive approach

Although we are not aware of any papers that do so, it would be straightforward to develop a model of world trade in which comparative advantage varied smoothly across space.

Imagine, for example, a Ricardian model in which a continuum of locations is arrayed in a line from north to south, and in which climate and therefore the relative productivity of labor in wine as opposed to wheat production varies smoothly with latitude. Then one would immediately have a model in which equilibrium in the world economy could be thought of in terms of the boundary between wheat and wine areas rather than in terms of the specialization of nations. It would be natural to think of this world as consisting of two regions, one producing wheat and the other wine, but the boundary between these regions would be endogenous rather than specified in advance. For some purposes this approach to modeling a seamless world might well prove useful; comparative advantage still explains much, perhaps most of world trade. However, both traditional location theory and recent work in economic geography generally assume away inherent differences between locations, and instead explain regional specialization in terms of some kind of external economies. We will remain within that tradition, and thus try to model a world in which economic differentiation among locations is entirely spontaneous.

In Krugman and Venables (forthcoming) we introduced a particular way of modeling external economies that seemed particularly well suited to thinking about geographical issues. We assumed that in each location labor must be allocated between two sectors: agriculture, where labor faces decreasing returns (because

of limited land supplies), and a manufacturing sector consisting of many differentiated products, each produced with increasing returns. In order to create external economies, we assumed that manufactured goods are used as intermediate inputs into each others' production. In the context of a two-location model, we were able to show that this setup could produce a spontaneous differentiation of roles between the locations, with one concentrating on manufacturing while the other concentrated on agriculture. In this paper we adopt the same general framework, but now imagine a world with a large number (in the simulations) or continuum (in the algebra) of locations. In this type of model it no longer makes sense to identify locations with regions; rather, we find that the model implies that groups of nearby locations will concentrate on manufacturing or agriculture, and we call these fuzzy-edged zones of economic orientation manufacturing or agricultural regions.

Imagine, then, a world in which both labor and arable land are evenly divided among many locations. Labor in each location can be used either to work the land or to produce manufactured goods; as in our previous model, we imagine that manufacturing is subject to increasing returns at the level of the firm and, because of the use of manufactured goods as inputs into each others' production, at the level of the industry. We also suppose that there are costs of transporting manufactured goods. In such a world, is

there any reason why the fraction of labor employed in manufacturing might vary across locations?

There is clearly an equilibrium in which all locations are identical -- that is, in which the share of the work force employed in manufacturing is the same everywhere. Let us call this the Flat Earth equilibrium. While the Flat Earth equilibrium always exists, however, it may not be stable. To see why, imagine perturbing it so that the initial share of labor in manufacturing varies slightly across the landscape, and consider a location whose neighborhood (in some sense) is somewhat more oriented toward manufacturing than the average. Such a location may well then turn out to be a place in which manufacturing is especially profitable, for two reasons. First, because manufactured goods are an input into manufactures production, a location that is surrounded by manufacturing-oriented locations will have access to a bigger market than one that is not; this "backward linkage" will tend to make industrial production there profitable. And at the same time, the availability of the inputs provided by nearby manufacturers will also encourage local industrial production -- an effect which corresponds to the idea of "forward linkage".

These effects will apply in reverse to a location whose neighborhood is less manufacturing-oriented than average: the lack of good access either to markets or to the inputs produced by other manufacturers will tend to make it an unattractive place to engage

in manufacturing.

So manufacturing will tend to be an attractive activity in a neighborhood that is already manufacturing-oriented, an unattractive activity in one that is not. If manufacturing expands where it is profitable and contracts where it is unprofitable, these differences among neighborhoods will tend to grow over time. That is, the world will spontaneously organize itself into manufacturing regions -- neighborhoods with a large share of the labor force employed in manufactures -- and agricultural regions -- neighborhoods with a small share so employed.

Can we say anything about how big manufacturing regions will be, and how they will be spaced? Loosely, we might argue that there will tend to be a typical size of and separation between such regions. If a manufacturing region is too big, its agricultural hinterland will be too large (that is, some of the farmers will be too distant to be a useful market); if a manufacturing region is too small, the forward and backward linkages will be insufficient to feed its growth. And of course if two manufacturing regions are too close together, they will either tend to merge or will inhibit each other by competing for the same hinterland. So while it may be wholly a matter of historical accident which regions become industrialized, the size and spacing of these regions should be at least roughly determined by the underlying parameters of the economy.

We can, however, be much more specific than this if we adopt an approach suggested by Alan Turing, in his landmark 1952 paper on morphogenesis.

Turing's paper was concerned with the process by which a group of initially identical cells become differentiated into the different parts of a developing embryo. The specific chemical mechanisms he described are still unsubstantiated in their original context, and in any case need not concern us here. What we can make use of is the remarkable way in which Turing simplified the analysis of a nonlinear system that seems at first sight hopelessly intractable.

The Turing approach involves two main simplifications. First, we need to impose a specific "geometry" on our spatial system. Turing assumed that his model embryo consisted of a ring of cells -- that is, a one-dimensional system with no ends. In our context this means assuming that the world is a "racetrack" economy, in which resources are distributed evenly around a circle. Second, Turing focussed on the initiation of cell differentiation rather than its conclusion; in our case, this means that we look at the early stages in which the world evolves away from the Flat Earth equilibrium, rather than at the long run equilibrium.

Suppose, then, that we imagine a world consisting of N locations (where N may be as large as one likes), arrayed in a circle (so that location N is next to location 1). Let m_j be the share of

manufacturing in location j 's employment. The implication of the story we have just told is that the rate of change of each m_j is a function of the vector of locational employment shares (and only of that vector, since it fully describes the state of the world at any instant). The details of that function depend on the whole structure of the model; but we can simply imagine a reduced form equation for the evolution of, say, m_1 of the form

$$\dot{m}_1 = F(m_1, m_2, \dots, m_N) \quad (1)$$

What is the equation for the evolution of m_2 ? Well, it involves the same function, but with the role of each location moved one step to the left:

$$\dot{m}_2 = F(m_2, m_3, \dots, m_N, m_1) \quad (2)$$

and similarly around the circle.

Now linearize this dynamic system around the Flat Earth equilibrium. Let m' be the vector of deviations of manufacturing employment from their Flat Earth values m^F ; then the dynamics of the model in the vicinity of the Flat Earth may be described as a system of the form

$$\dot{m}' = Am' \quad (3)$$

where A is an N by N matrix.

The way one always analyzes such a linear system is to look

for the eigenvectors of A and their associated eigenvalues. Turing, however, pointed out that A has a special structure that simplifies this task. Consider (1) and (2); by linearizing these two equations around $m_j = m^F$ we get the first two rows of A . But that means that the second row of A is the same as the first row, but displaced one location to the right (and with the last element wrapped around to the front). Generally then, the element of A in row r and column s can be written a_{s-r} , where it is understood that if the index is negative, N is added to its value. In other words, A has a cyclic structure.

This structure allows us to construct the eigenvectors very easily. Suppose that

$$m_r = e^{i\phi r} \quad (4)$$

for all r , where i is the square root of minus one and we assume that $\phi N/2\pi$ is an integer. This is an eigenvector if it satisfies the following equations, for each $r = 1..N$,

$$\sum_{s=1}^N a_{s-r} e^{i\phi s} = \lambda e^{i\phi r} \quad (5)$$

where λ is the eigenvalue. Rearranging,

$$\sum_{s-r=1}^N a_{s-r} e^{i\phi(s-r)} = \lambda \quad (6)$$

This is the same equation for all rows r , so (4) indeed defines an eigenvector; for each value of ϕ the eigenvalue is given by (6).

The eigenvector defined in (4) has elements which are complex rather than real numbers. However, since

$$e^{ix} + e^{-ix} = 2\cos(x) \quad (7)$$

and

$$i(e^{ix} - e^{-ix}) = 2\sin(x) \quad (8)$$

these complex eigenvectors can be reassembled into simple sine and cosine waves. In other words, any sinusoidal fluctuation in manufacturing employment whose wavelength is an integer fraction of the circumference of the economy will be an eigenvector, with an eigenvalue that depends only on that wavelength.

But such sinusoidal fluctuations in manufacturing have another aspect: one can represent the vector m' as a Fourier series, that is, as a sum of sine and cosine waves with wavelengths N , $N/2$, $N/3$, and so on. What we have just demonstrated is that the components of that Fourier representation are also eigenvectors. Thus the evolution of the economy near the Flat Earth equilibrium can in effect be represented as the parallel and independent evolution of a set of shadow economies, in each of which the variation of manufacturing employment is a perfectly regular sine wave.

We can now think of the process described loosely above, in which slight initial variations in manufacturing orientation lead the economy to organize itself into manufacturing and agricultural regions, more precisely: such self-organization will take place if and only if at least one of the sinusoidal eigenvectors of the dynamic system has a positive eigenvalue. If it does, then fluctuations of that wavelength will grow over time, producing regional differentiation.

What if the eigenvalues of more than one wavelength fluctuation are positive? If we start close enough to the Flat Earth, by the time the unevenness of manufacturing distribution becomes marked it will be dominated by whichever fluctuation has grown most rapidly -- that is, by the wavelength with the largest eigenvalue. The answer to the question, how far apart are manufacturing regions likely to be, is therefore answered by finding this "preferred" wavelength.

Notice, by the way, the role of initial conditions in this analysis. The preferred wavelength is determined by the parameters of the model; thus the number and spacing of manufacturing regions will not depend on initial conditions, as long as the initial distribution is sufficiently flat. The initial distribution of manufacturing will, however, determine the phase of the fluctuation at that preferred wavelength, and thus which specific locations become industrialized.

It is possible, then, to think about the emergence of a

spatially structured world economy in terms of the growth of "Turing waves", regular fluctuations across space in the manufacturing share of employment. While this approach offers a startlingly simple way of thinking about a seemingly complex issue, however, it has two important limitations.

First, even if we consider the evolution of a world economy that starts from a close approximation to the Flat Earth, Turing wave analysis applies only to the initial period of self-organization. Later on, as the distribution of manufacturing becomes very uneven, the nonlinearity of the dynamic system becomes unavoidable; so Turing analysis cannot give us a good picture of what happens thereafter. To analyze these later stages we must turn to numerical methods.

Second, while it is useful to carry out the exercise of asking how an economy evolves a spatial structure by diverging away from the Flat Earth, this is clearly what we might call a fictitious history. Perhaps it is not a bad description of, say, the emergence of Flanders and northern Italy as manufacturing centers in the early Middle Ages. But from then on each successive spatial structure emerged from an economy that was already structured into industrial and farming regions, so that once again Turing analysis is no longer valid.

To some extent this means that the dynamics of economic geography fit Henry Ford's definition of history, as "just one damn

thing after another". We may, however, be able to get some more insight by imagining a different, slightly more realistic fictitious history. In this history we start with an economy that already has a regular structure of manufacturing and agricultural regions. We then gradually change its parameters -- in particular, it is interesting to explore what happens if transport costs fall, leading to growing world integration. What one expects to happen in this case is an eventual bifurcation: at some point the initial structure becomes unstable, and the economy makes a transition to a new structure, with a different (presumably smaller) number of manufacturing regions. If one then continues the process, this structure in turn eventually collapses, and so on. While this more complex fictitious history cannot be treated using Turing's simplification and can thus be explored only via simulations, by following only a particular path through the many possible equilibria of the system we may be able to avoid wandering off into a taxonomic swamp.

This is about as far as we can get without being more specific about the details. We therefore turn next to the statement and solution of a spatial model of the world economy.

3. A formal model

The structure of this model is closely related to that in Krugman and Venables (forthcoming), and makes use of the same sorts of technical tricks, involving the combination of Dixit-Stiglitz

monopolistic competition and iceberg transport cost, that have figured in many recent papers in economic geography. We will therefore be brief in describing its formal structure.

We consider a world in which economic activity takes place on the circumference of a circle whose radius we normalize to π . Two immobile factors -- workers and land -- are uniformly distributed around that circumference with densities L and K respectively. Each location can produce both manufacturing and agricultural output.

Manufacturing

Manufacturing is assumed to be a monopolistically competitive sector, in which the number (or more formally the density) of firms operating at location r is endogenously determined, and denoted $n(r)$. The producer price charged by firms at r is $p(r)$. Shipments of manufactured goods are subject to "iceberg" transportation costs -- that is, a fraction of any shipment melts away in transit. The number of units that must be shipped from r in order that one unit arrives at s is denoted $t(r,s)$ and given by

$$t(r,s) = e^{tD(r,s)} \quad (9)$$

Trade costs are therefore an exponential function of distance, where $D(r,s)$ is the (shorter!) distance from r to s around the

circumference of the circle.

We will assume that different manufactured goods can be aggregated via a CES function to yield a manufactures composite that is used both as a consumption good and as an intermediate input. This CES function may be represented indirectly by a CES price index, q . In each location the price index is defined over products supplied from all sources; at location r it takes the form:

$$q(r)^{1-\sigma} = B \int_{r-\pi}^{r+\pi} n(s) [p(s)t(s,r)]^{1-\sigma} ds \quad (10)$$

where B is a constant, and σ (>1) is a measure of product differentiation. Demand for each variety is derived by partial differentiation of this price index with respect to the price of the variety. If expenditure on manufacturing in location s is denoted $c(s)$, then demand in s for a single variety produced in r and sold in s , $x(r,s)$, is:

$$x(r,s) = B p(r)^{-\sigma} t(r,s)^{1-\sigma} q(s)^{\sigma-1} c(s). \quad (11)$$

Integrating around the circle, the total demand for a product produced in r , $x(r)$, is,

$$x(r) = p(r)^{-\sigma} B \int_{r-\pi}^{r+\pi} t(r,s)^{1-\sigma} q(s)^{\sigma-1} c(s) ds. \quad (12)$$

A single firm in location r has profits given by:

$$\Pi(r) = p(r)x(r) - [(1-\mu)w(r) + \mu q(r)] [\alpha + \beta x(r)]. \quad (13)$$

The first term is revenue, and the second costs. We assume a fixed input requirement of α and a constant marginal input requirement β . The input is a fixed coefficient aggregate of labour and intermediate goods, with input requirements of labor and intermediates $1-\mu$ and μ respectively; $w(r)$ is the manufacturing wage rate, and $q(r)$ the price index of intermediates. In other words, each firm uses both labor and all varieties of industrial products, where we assume that these are appropriately aggregated by the CES form of equation (10).

Since the producer of an individual good faces an elasticity of demand σ , firms mark up price over marginal cost by the factor $1/\sigma$. We choose units of measurement such that $\beta\sigma = \sigma - 1$, so that price is

$$p(r) = (1-\mu)w(r) + \mu q(r). \quad (14)$$

Firms are scaled such that they earn zero profits at size 1; this is done by setting $\alpha = 1/\sigma$. We assume instantaneous entry and exit of firms, so that

$$x(r) = 1. \quad (15)$$

Operating at this scale, each firm demands $(1-\mu)$ units of labor. Manufacturing employment at location r is therefore

$$m(r) = (1-\mu)n(r). \quad (16)$$

Agriculture and general equilibrium.

Agricultural output, which we choose as numeraire, is assumed costlessly tradeable. At each location the agricultural production function is Cobb-Douglas in land and labour, with labour share η . Workers are geographically immobile but sectorally mobile, so the agricultural wage, $w^A(r)$ is given as the (value) marginal product of labour in agriculture,

$$w^A(r) = \eta K^{1-\eta} [L-m(r)]^{\eta-1}. \quad (17)$$

Combining intermediate and consumer demands we can find expenditure on manufacturing,

$$c(r) = \mu n(r) q(r) + \gamma [(1-\mu)n(r)w(r) + K^{1-\eta}(L-m(r))^\eta]. \quad (18)$$

The first term is intermediate demand and the second consumer expenditure. Consumers devote a proportion γ of their income to expenditure on differentiated products, and total income is given in the square brackets; within the brackets the first term is income from manufacturing employment and the second is total income generated in agriculture.

Instantaneous equilibrium.

Rather than supposing that the economy moves immediately to a full equilibrium, we introduce dynamics by supposing that there is a gradual process of reallocation of labor between manufacturing and agriculture. Specifically, we assume that

$$\dot{m}(r) = \delta[w(r) - w^A(r)]. \quad (19)$$

At any instant, then, we think of the economy as having a predetermined allocation of labor to manufacturing at each location, $m(r)$. Corresponding to this allocation is a short-run equilibrium. For given values of $m(r)$ equations (10), (12), (14), (15), (16), (17), and (18) give equilibrium values of $q(r)$, $p(r)$, $x(r)$, $w(r)$, $n(r)$, $w^A(r)$ and $c(r)$.

If we define long run equilibrium as a stationary state in which the allocation of labor is no longer changing, then such a long-run equilibrium requires that manufacturing and agricultural wages be equal. One such long run equilibrium is the Flat Earth case. We choose units such that in this case,

$$\begin{aligned} m(r) &= q(r) = w(r) = p(r) = 1, \\ n(r) &= c(r) = 1/(1-\mu). \end{aligned} \quad (20)$$

This can be achieved by appropriate choice of constants B , L and K (values are given in the appendix).

Our next step is to explore the process by which this

economy diverges away from the Flat Earth equilibrium.

3. The emergence of spatial structure

One way to explore the dynamics of this model is simply to assume a set of parameters, impose an initial allocation of labor between agriculture and manufactures, and run the model on a computer. (It is of course necessary for simulation purposes to work with a number of discrete locations rather than continuous space).

Figure 1 shows the results of a typical run. The figure plots the share of manufacturing in employment at each location (on the vertical axis) as a function of time. We consider a world with 40 locations, (bear in mind that location 40 is next to location 1). The starting point for this run was a position very close to the Flat Earth equilibrium. Specifically, we set $m(r)$ equal to one plus 0.0001 times random variables drawn from an $N(0,1)$ distribution; parameters take values, $\mu = 0.5$, $\sigma = 6$, $\eta = 0.9$, $\gamma = 0.3$ and $\tau = 6$, where $\tau = e^{\pi}$, i.e., is the number of units that have to be shipped to secure delivery of one unit to the most distant location.

Although the world economy in this run starts off with almost no spatial structure, we see that it eventually evolves a structure with a strong differentiation between industrial and agricultural regions. This emergent structure is also, it turns out, extremely regular: the three industrial regions are evenly spaced around the circle. Repeated runs of this model, with initial

conditions randomly chosen each time, consistently produce this result.

Our discussion in part 1 of this paper already suggests why. What is happening in Figure 1 is that the divergence away from the Flat Earth is being driven by a Turing wave of frequency 3; by the time a linearized model is no longer a good approximation, the pattern of peaks and troughs in manufacturing orientation has become sufficiently pronounced to lock in the eventual position of industrial and farming regions.

We can further confirm this interpretation of the dynamics by showing how the predicted rate of growth of a Turing wave depends on its frequency. Consider the model described by equations (10) - (19). If we linearize this model around the Flat Earth equilibrium, we get the following system:

$$\begin{aligned}
(1-\sigma)q' &= \frac{B}{1-\mu} \int_{r-\pi}^{r+\pi} [(1-\mu)n' + (1-\sigma)p'] \psi(s,r)^{1-\sigma} ds, \\
\sigma p' &= \frac{B}{1-\mu} \int_{r-\pi}^{r+\pi} [(\sigma-1)q' + (1-\mu)c'] \psi(s,r)^{1-\sigma} ds, \\
p' &= (1-\mu)w' + \mu q', \\
m' &= (1-\mu)n', \\
(w^A)' &= \zeta m', \quad \zeta \equiv \frac{1-\eta}{\eta} \frac{\gamma}{1-\gamma}, \\
c' &= [\mu + \gamma(1-\mu)]n' + \frac{\mu}{1-\mu} q' + \gamma w' - \gamma m'.
\end{aligned}
\tag{21}$$

where ζ is the elasticity of agricultural wages with respect to industrial employment. We simplify these equations in two ways. First, eliminate endogenous variables p' , c' , and n' . Second, define constants α_w and α_q (the values of which remain to be found), such that

$$w'(r) = \alpha_w m'(r), \quad q'(r) = \alpha_q m'(r).
\tag{22}$$

Equations (21) can then be expressed as:

$$\begin{aligned}
(1-\sigma)\alpha_q &= \left[1 + (1-\sigma)((1-\mu)\alpha_w + \mu\alpha_q)\right] \\
&\quad \cdot \frac{B}{1-\mu} \int_{r-\pi}^{r+\pi} m'(z) t(s,r)^{1-\sigma} dr \\
\sigma[(1-\mu)\alpha_w + \mu\alpha_q] &= \left[\mu + \gamma(1-\mu)\alpha_w + (\mu+\sigma-1)\alpha_q\right] \\
&\quad \cdot \frac{B}{1-\mu} \int_{r-\pi}^{r+\pi} m'(z) t(s,r)^{1-\sigma} ds.
\end{aligned}
\tag{23}$$

The differential equation for m' , equation (19), takes the form,

$$\dot{m}' = \delta[\alpha_w - \zeta]m'. \tag{24}$$

these equations do not look very friendly. We know from part 1, however, that we need concern ourselves only with sinusoidal fluctuations in $m'(r)$. Suppose, then, that we assume that

$$m'(r) = \cos(\phi r) \tag{25}$$

with ϕ an integer. This enables us to evaluate the definite integrals in (23);

$$\begin{aligned}
\frac{B}{1-\mu} \int_{r-\pi}^{r+\pi} m'(s) t(s,r)^{1-\sigma} dz &= \frac{B}{1-\mu} \int_{r-\pi}^{r+\pi} \cos(\phi s) (e^{t|s-r|})^{1-\sigma} ds \\
&= H(\phi) \cos(\phi r)
\end{aligned}
\tag{26}$$

where

$$H(\phi) = \frac{t^2(\sigma-1)^2}{t^2(\sigma-1)^2 + \phi^2} \left[\frac{1 + e^{i(1-\sigma)\pi}(-1)^\phi}{1 - e^{i(1-\sigma)\pi}} \right] \quad (27)$$

Values of α_w and α_q can now be obtained as functions of $H(\phi)$ from equations (23), and an explicit equation for α_w is given in the appendix.

We saw in equation (24) that the rate of growth of a fluctuation of frequency ϕ -- which is to say the eigenvalue of that eigenvector -- is proportional to $[\alpha_w - \zeta]$. This term can be used to find the "preferred" frequency.

Figure 2 shows α_w as a function of ϕ for the parameters used to generate Figure 1. α_w is positive for frequencies 2 and above, and approaches zero from above as the frequency goes to infinity (this property is quite general). α_w captures the fact that locations with more manufacturing pay a higher manufacturing wage. Against this, locations with more industry have a higher agricultural wage, and this effect is captured by the parameter ζ , and illustrated by the line $\zeta\zeta$. The growth rate is positive for fluctuations of frequencies 2 - 6; thus this economy can be expected to diverge away from the Flat Earth. And it shows that the growth

rate is highest for fluctuations of frequency 3, thus confirming that the emergence of a regular spatial structure in this case can be interpreted as the growth of a Turing wave.

We might ask why an intermediate wavelength has the highest growth rate -- which is equivalent to asking why the economy develops several distinct industrial regions, as opposed to a single center or a wide dispersion. The basic intuition was already given in part 1: too large an industrial region (too long a wavelength) will not have enough close-in hinterland, while too small a region (too short a wavelength) will not generate sufficient linkages.

The number of industrial regions that emerge depend in a systematic way on the parameters of the model. Table 1 shows how the preferred frequency depends on the share of intermediates in production and the level of transport costs (units shipped for delivery of one unit at distance π). The integer values in the table give the preferred frequencies, and we see that strong linkages (high μ) and low transport costs give a lower preferred frequency.

The bracketed numbers in the table give the value of α_w at the preferred frequency. Concentrations form if this number exceeds ζ . If transfer of labor from agriculture to industry has very little effect on the agricultural wage then even weak linkages will cause agglomeration. However, when agricultural wages are more elastic with respect to manufacturing employment then substantial

linkages are needed for the eigenvalue to be positive; (recall that $\zeta = 0.0476$ in the examples of figures 1 and 2). Agglomeration therefore requires strong linkages and/or low ζ . When agglomeration occurs, strong linkages and low transport costs give fewer, larger concentrations.¹

Table 1: Preferred Frequencies, $(\alpha_w(\phi))$.

	$\mu = 0.1$	$\mu = 0.3$	$\mu = 0.5$
$\tau = 1.5$	2 (0.0022)	1 (0.025)	1 (0.076)
$\tau = 3$	5 (0.0022)	3 (0.024)	2 (0.092)
$\tau = 6$	8 (0.0022)	4 (0.025)	3 (0.094)
$\tau = 20$	14 (0.0022)	7 (0.025)	5 (0.094)

4. The evolution of spatial structure

Up to this point we have considered only how a Flat Earth might organize itself into agricultural and manufacturing regions.

¹ Reducing σ (raising price cost mark-ups) also has the effect of leading to fewer and larger concentrations.

Ordinarily, however, it is more interesting to ask how changes in underlying parameters will alter the pattern of trade and specialization in a world that is already differentiated into such regions. In particular, what will happen to the pattern of specialization as the world economy becomes increasingly well integrated over time?

We may address this question by carrying out the following experiment. We reduce the transportation cost t in a series of small steps. Following each step, we allow the model economy to evolve until it reaches a steady state; then take the next step.

Anyone who is familiar with nonlinear systems can guess what might happen in such an experiment. As t falls -- in effect, as the world gets smaller -- the model eventually reaches a point at which an equilibrium structure of manufacturing regions is no longer stable. At that point, then, this structure will unravel, giving rise to a new structure with fewer manufacturing regions. As we continue to reduce transport costs, this new structure will persist for a time; then it in turn will collapse, and so on.

And that is indeed exactly what we find. Figure 3 summarises the results of such an experiment (constructed with the same parameter values as figure 1). The bold solid lines give the "preferred" frequencies that will be reached starting from the Flat Earth equilibrium. The dashed lines and downwards pointing arrows show the effects of starting with high transport costs and a

four manufacturing region structure, and progressively reducing transport costs.

As transport costs fall, there is a critical point at which the four-manufacturing-region structure collapses into a three-region structure. This then collapses into a two region structure and then into a world with a single manufacturing region. We find, then, that even if the change in underlying parameters is gradual, the evolution of the world's spatial structure is characterized by "punctuated equilibrium". Long stretches of stability are interrupted by episodes of discontinuous change.

We should also note that the equilibria that the model "visits" as we reduce transport costs do not coincide with the equilibria attained starting from a Flat Earth. There is therefore a path dependence in the structure of equilibria. This is emphasised further if we run history backwards: start with the one-manufacturing-region equilibrium and gradually ~~increase~~ τ . This is illustrated by the dashed lines and upwards pointing arrows. The model retraces its steps, from one to two and then to three manufacturing regions, but transitions do not occur at the same levels of τ ; clearly these ranges are overlapping.

We see, then, that this model exhibits a considerable degree of path dependence. Even in the story of self-organization from a near-uniform world, which locations become industrial depends on details of the initial conditions. But beyond this, as the economy

evolves over time, even gross features of the spatial structure of production may depend on where the economy has come from.

Figure 4 gives information about the phase of the manufacturing regions that develop as transport costs are reduced. The two-manufacturing-region world develops from the three-region world, with one region at the same location, and the other midway between two of the previous manufacturing regions. The one-manufacturing-region world has its manufacturing centre midway between, although intersecting with, the manufacturing regions of the two region world.

5. Conclusions

The idea of modeling international trade simply as flows across a landscape, without the usual assumption that countries or at least regions can be represented as points, at first seems both outlandish and likely to be analytically intractable. We hope, however, that we have succeeded in making the case that thinking in terms of a world economy in which activity is spread continuously across space is perfectly reasonable, at least as an intellectual exercise. Indeed, once one has worked with a spatial model of the global economy, the traditional assumption that countries are dimensionless points can be seen as the awkward if useful assumption it is, rather than as a particularly compelling way to think about international specialization and trade. And it turns out

that such a spatial model need not be analytically intractable; with a little help from the computer, and a lot of help from ideas created in other fields, it is possible to develop some quite clear intuitions about how spatial structures emerge and evolve in a "seamless" world.

In addition to suggesting an alternative way to think about world trade, we hope that this paper has served as something of an advertisement for the use in international economics of concepts and techniques borrowed from fields outside economics. In the model developed here, such typical features of nonlinear systems as path dependence and punctuated equilibria occur in a very natural way; while these concepts have received some application in economics, they are still used only occasionally, and it seems useful to show that they may be of direct relevance to the understanding of as basic a subject as international specialization and trade. And we are of course delighted to have found a compelling economic application of Turing's beautiful analysis of the origins of spatial structure.

It remains true, however, that the framework that we have used to describe world trade is aggressively unrealistic. Thus one would like to know in which directions it should be extended to move it somewhat closer to applicability.

One extension which sounds simple is to go from the circular "racetrack" economy to a two-dimensional surface. Preliminary work suggests that in practice this extension is anything

but simple: two-dimensional spatial models quickly become astonishingly difficult both to analyze with paper and pencil and even to simulate on a computer (the time needed to do a simulation run can easily go from minutes to days). Nonetheless, like it or not we do not live in a one-dimensional economy, and it is important to know what difference that makes.

Another extension would be to add some nonuniformity in the landscape. It is common practice among historians to argue that aspects of natural geography, such as the availability of navigable rivers, have played a crucial role in the differential growth of different regions; this framework offers one way to model such “catalytic” effects formally.

Finally, it would be natural to extend the framework to multiple industries, so that what emerges is not simply a division of the world into agricultural and manufacturing regions, but specialization within the manufacturing sector.

It is easy to see that the initial simplicity of our approach, especially the story of how spatial order emerges from a Flat Earth economy, can quickly dissipate as we add realistic complications. Nonetheless, we hope that this model has demonstrated that a truly spatial model of world trade is both possible and worth doing.

Appendix

The Flat Earth equilibrium has endogenous variables taking the values given in (20) if units are chosen such that parameters B, L and K satisfy:

$$\begin{aligned}\frac{1-\mu}{B} &= \int_{-\pi}^{\pi} t(z,y)^{1-\sigma} dz = 2 \left[\frac{e^{\eta(1-\sigma)\pi} - 1}{t(1-\sigma)} \right], \\ L &= 1 + \eta \left(\frac{1-\gamma}{\gamma} \right), \\ K &= \eta^{\eta/(\eta-1)} \frac{1-\gamma}{\gamma}.\end{aligned}\tag{28}$$

α_w can be expressed as a function of $H(\phi)$ by solving equations (23) to give:

$$\begin{aligned}\alpha_w &= \frac{1}{\Delta(\sigma-1)} \left[\frac{\mu(1-2\sigma)}{H(\phi)} + \mu^2(\sigma-1) + \mu + \sigma - 1 \right] \\ \Delta &= (\mu-1) \left[\frac{\sigma}{H(\phi)^2} - \frac{\gamma}{H(\phi)} + \mu\gamma + 1 - \sigma - \mu \right]\end{aligned}\tag{29}$$

As $\phi \rightarrow \infty$, $H(\phi) \rightarrow 0$, and $\alpha_w \rightarrow 0$ from above, as illustrated in figure 2.

REFERENCES

- Deardorff, Alan V. and P.N. Courant, (1992),
"International trade with lumpy countries",
Journal of Political Economy, 100(1), February
1992, 198-210.
- Dixit, A. and Stiglitz, J. (1977), "Monopolistic competition and
optimum product diversity", *American Economic
Review*.
- Krugman, P. (1991), "Increasing returns and economic
geography", *Journal of Political Economy*, June.
- Krugman, P. and A.J. Venables, (1995), "Globalization and the
inequality of nations", *Quarterly Journal of Economics*,
forthcoming.
- Krugman, P. (forthcoming), *The Self-Organizing Economy*,
Cambridge Mass.: Blackwell.
- Turing, Alan M. (1952), "The chemical basis of
morphogenesis", *Philosophical Transactions of the Royal
Society*, 237, 5-72.
- Venables, A. J. (1993), "Equilibrium locations of vertically
linked industries", CEPR discussion paper no. 802.

Figure 1




Figure 2




Figure 3




Figure 4: Global Transition

