Internet cluster emergence

by

Danny Quah *
LSE Economics Department
January 2000

*I thank the ESRC and the Andrew Mellon Foundation for financial support. Larry Christiano, Jacques Crémer, Bob Lucas, Peter Neary, and seminar participants at the LSE, Northwestern, the University of Sydney, and a MacArthur Foundation "Social interactions and economic inequality" group meeting provided helpful comments on related work.

Nontechnical Summary

Why does economic activity locate where it does? Knowing this informs discussion of national productivity and competitiveness. It aids understanding the implications of policies that are regionally- or nationally-based.

Economic geography provides a now-standard answer: Location reflects a tradeoff between costly transportation on the one hand and increasing returns or some other positive externality on the other. The latter would lead to all economic activity locating in one place alone, while the former to economic production diffusing outwards for conveniently transporting output to users (whether final consumers or other producers). Equilibrium balances the two tensions, and produces nondegenerate landscapes of economic activity.

Internet economic activity poses a challenge to this conventional wisdom in its promise to dramatically reduce transportation costs. Of course, such a promise likely changes only marginally decisions in, say, the manufacturing of construction cranes or heavy machinery, or in the mining of petroleum. Sectors of the economy most affected would be where output is intangible or "weightless". But examples of those are legion, and include some of the fastest-growing industries in any modern developed economy: financial and consulting services, software, health consulting, music and entertainment, and similar others.

The puzzle, from the perspective of standard theory, is that these industries show neither all activity occurring in just a single location nor, the opposite, a completely random scattering over geographical space. Instead, distinctive of industries like finance or software development is their clustering into specific locations spaced at almost regular intervals about the globe. Bangalore has become a powerhouse of software codecutting. It generated a high proportion of India's US$3 billion 1999 software revenue (at 50% growth per year since 1992, and of which 60% have been export earnings). Even better known are Finland, Singapore, Ireland, and Silicon Valley USA as centers of advanced information technology development. By the same token, Tokyo, London, and New York constitute prominent clusters of

financial activity.

Of course, any number of ready explanations might account for some of these observations. Universities and research centers attract other knowledge-intensive activities to cluster around them. A millenium of history in banking and stockbroking, and plain old inertia explain the location of some prominent financial centers. But not all locations of higher learning are also successful centers of these related activities. Nor do all such clusters locate around obvious, already-extant centers of research and learning. Indeed, some of these observations simply beg the deeper question of why knowledge and intellectual input are geographically localized at all since, a priori, they shouldn't be.

This paper takes a different approach to explaining observations on clustering. It develops a model where transportation costs don't matter and where the underlying geography is homogenous. Equilibria emerge that, nonetheless, show distinct clusters in the location of economic activity. Thus, the distribution of observable outcomes ends up more skewed or unequal than the distribution of underlying characteristics—an effect related to the economics of superstars. (Geographers are not unambiguous on what a cluster is. In the current work, the connection with the economics of superstars could well be taken to *define* the spatial clusters of interest.)

In the model, agents optimize dynamically and have rational expectations over outcomes. Equilibrium is a law of motion in spatial distributions over a 3-dimensional globe. Clusters appear as waves in time and space. They arise along convergent transition paths—in the space of distributions—to long-run steady state. These periodic waveforms result, in the model, from a tradeoff between productivity spillovers across timezones and a transient "stickiness" in factor input location.

Internet cluster emergence

by

Danny Quah
LSE Economics Department

January 2000

ABSTRACT

Internet development holds the promise of transmitting economic value across physical space at zero marginal cost. In such a "weightless economy", what factors matter for the location of economic activity and thus for economic development? This paper sketches a model of spatial dynamics over a three-dimensional globe, where transportation costs don't matter. The paper develops conditions under which clusters of activity emerge.

Keywords: distribution dynamics, economic geography, Internet, location, space, time, Toeplitz, waves, weightless economy

JEL Classification: D30, O10, O18, O33

Communications to: D. Quah, LSE, Houghton Street, London WC2A 2AE.

Tel: +44 207 955-7535, Email: `dquah@econ.lse.ac.uk`
(URL) `http://econ.lse.ac.uk/~dquah/`

1 Introduction

Why does economic activity locate where it does? Economic geography provides a now-standard answer: Location reflects a tradeoff between costly transportation on the one hand and increasing returns or some other positive externality on the other. The latter would lead to all economic activity locating in one place alone, while the former to economic production diffusing outwards for conveniently transporting output to users (whether final consumers or other producers). Equilibrium balances the two tensions, and produces nondegenerate landscapes of economic activity.[1]

Internet economic activity poses a challenge to this conventional wisdom in its promise to dramatically reduce transportation costs. Of course, such a promise likely changes only marginally decisions in, say, the manufacturing of construction cranes or heavy machinery, or in the mining of petroleum. Sectors of the economy most affected would be where output is intangible or "weightless". But examples of those are legion, and include some of the fastest-growing industries in any modern developed economy: financial and consulting services, software, health consulting, music and entertainment, and similar others.

The puzzle, from the perspective of standard theory, is that these industries show neither all activity occurring in just a single location nor, the opposite, a completely random scattering over geographical space. Instead, distinctive of industries like finance or software development is their clustering into specific locations spaced at almost regular intervals about the globe. Bangalore has become a powerhouse of software codecutting. It generated a high proportion of India's US$3 billion 1999 software revenue (at 50% growth per year since 1992, and of which 60% have been export earnings). Even better known are Finland, Singapore, Ireland, and Silicon Valley USA as centers of advanced information technology development. By the same token, Tokyo, London, and New York constitute prominent clusters of

[1] Examples of such reasoning are in, e.g., Fujita, Krugman and Mori (1999) and Krugman and Venables (1997).

financial activity.

Of course, any number of ready explanations might account for some of these observations. Universities and research centers attract other knowledge-intensive activities to cluster around them. A millennium of history in banking and stockbroking, and plain old inertia explain the location of some prominent financial centers. But not all locations of higher learning are also successful centers of these related activities. Nor do all such clusters locate around obvious, already-extant centers of research and learning. Indeed, some of these observations simply beg the deeper question of why knowledge and intellectual input are geographically localized at all since, a priori, they shouldn't be (Arrow, 1962). But this, of course, turns out to be the same question with which we began above.

This paper takes a different approach to explaining observations on clustering. It develops a model where transportation costs don't matter and where the underlying geography is homogenous. Equilibria emerge that, nonetheless, show distinct clusters in the location of economic activity. Thus, the distribution of observable outcomes ends up more skewed or unequal than the distribution of underlying characteristics—an effect related to the economics of superstars[2] (Rosen, 1981).

In the model, agents optimize dynamically and have rational expectations over outcomes. Equilibrium is a law of motion in spatial distributions over a 3-dimensional globe. Clusters appear as waves in time and space. They arise along convergent transition paths—in the space of distributions—to long-run steady state. These periodic waveforms result, in the model, from a tradeoff between productivity spillovers across timezones and a transient "stickiness" in factor input location.

[2] Geographers are not unambiguous on what a cluster is. In the current work, the connection with the economics of superstars could well be taken to *define* the spatial clusters of interest.

2 Model: Cluster Emergence on the Globe

Consider the physical geography of a 3-dimensional globe, Fig. 1. When time matters, but not geographical distance, the globe is isomorphic to the Equator (provided we ignore the North and South Poles). View the Equator from the North Pole, and define geography G to be the unit circle in the complex plane $\{z \in \mathsf{C} : |z| = 1\}$ or, equivalently, $\{\omega : \omega \in (-\pi, \pi]\}$, with $z(\omega) = e^{i\omega}$ or $\omega(z) = i^{-1} \log z \pmod{\pi}$.

The model has two key features. First, productivity spillovers are characterized by *timeliness* or connections across time zones. Thus, longitudinal distance matters even though physical Euclidean metric does not. In the model, Helsinki is right next to Athens and only half the distance apart of London from Paris (whereas on a 3-dimensional globe, Helsinki/Athens is over 7 times the distance of London/Paris). Second, economic activity comprises forward-looking producers with rational expectations. These producers locate their factor inputs optimally on geography G by maximizing, over the future infinite horizon, a present discounted value of profit flows. The factor input is "sticky", i.e., its uprooting or its implanting uses up resources. But once uprooted, however far the input is moved entails no further costs: these are adjustment costs, not iceberg transportation costs.

2.1 Spillovers across the globe

For production at location z, what happens at z' matters through the latter's timeliness relative to events at z. Timeliness is a mapping $\mathsf{T} : \mathsf{G} \times \mathsf{G} \to [0, 1]$ such that (i) local production is always the most timely, i.e., for each z, the maximum value of 1 in $\mathsf{T}(\cdot, z)$ is attained at z, i.e., $\mathsf{T}(z, z) = 1 \geq \mathsf{T}(z', z)$ for all $z' \in \mathsf{G}$; and (ii) timeliness is invariant under equatorial rotation or is radially homogeneous: it varies with radial separation, not with where production occurs, i.e., for all z, z' in G, timeliness $\mathsf{T}(z', z)$ depends only on $\omega(z) - \omega(z') \pmod{\pi}$.

From radial homogeneity, for any positive real number r,

$$\int_{\mathsf{G}} |\mathsf{T}(z',z)|^r \, dz' = \int_{\mathsf{G}} |\mathsf{T}(z,z')|^r \, dz' = \int_{\mathsf{G}} |\mathsf{T}(z',1)|^r \, dz'.$$

We can therefore define

$$\|\mathsf{T}\|_r \stackrel{\text{def}}{=} \left(\int_{\mathsf{G}} |\mathsf{T}(z',1)|^r \, dz' \right)^{1/r} < \infty.$$

Assumption (i) does not disallow $\mathsf{T}(z',z) = 1$ at some other $z' \neq z$; it merely requires achieving the maximum at $z' = z$.

It is instructive to contrast timeliness with the more conventional physical distance metric. When spillovers are spatial, physical distance (or its inverse) provides a ready reckoning of the strength of spillover effects. Physical distance achieves its minimum at $z' = z$; it is symmetric in all directions, monotone increasing, and radially homogeneous. Timeliness (i) and (ii), while sharing the minimum and radial homogeneity properties, impose neither symmetry nor monotonicity in spillovers across G. The analysis to follow gives qualitatively the same implications regardless of whether "bumps" or cyclicalities appear in T. It is useful to allow these, perhaps to accommodate the sleep patterns of economic agents across different parts of the globe. But at the same time, because T *could* be asymmetric monotone decreasing, any cyclicalities emerging in equilibrium are not directly inherited from assumptions on T alone.

2.2 Producers

Denote by $f_t(z)$ the quantity of factor input located at z in time t. (Without ambiguity I will use the same symbol f also to refer to the factor input itself.) The total amount of factor input is constant through time, so that:

$$\int_{\mathsf{G}} f(z) \, dz = 1, \qquad f(z) \geq 0. \tag{1}$$

Equation (1) specifies f to be a spatial (probability) density across geography G. Maintaining (1), the function f will evolve through time.

A single homogeneous good Y is produced worldwide. At location z, output depends on the factor input not just at z but everywhere in G:

$$\forall z \in \mathsf{G}: \quad Y(z) = \left[\int_{\mathsf{G}} [\mathsf{T}(z', z) f(z')]^{\gamma}\, dz'\right]^{1/\gamma}, \qquad \gamma \in (0,1). \tag{2}$$

In (2), since f, in general, evolves through time, so does output Y. Throughout this paper we maintain T time-invariant.

The "spillover" effect of f at a particular location is maximized there, but f elsewhere contribute as well, with spillover strength depending on $\mathsf{T}(z', z)$. The (conditional) elasticity of substitution, $(\gamma - 1)^{-1}$, will be required below to be at least 1 in absolute value. Its critical range, however, depends on other parameters in the model, so we reserve the possibility of restricting γ further within $(0, 1)$.

Finally, uprooting or planting f is costly, although f's movement across space is not. Represent this by costs of adjustment:

$$C(\dot{f}_t(z)) = \frac{1}{2}\zeta \times \dot{f}_t(z)^2, \qquad \zeta > 0, \tag{3}$$

where $\dot{f}_t(z)$ denotes the time derivative of $f_t(z)$. The larger is ζ, the less easily does f change at a given point.

The producer at z decides how much f to have there by controlling $\dot{f}(z)$. He behaves competitively and earns on each unit of f in place a return W, equal to the marginal product of $f(z)$ in global production. To calculate this, first notice that from (2) the marginal product of $f(z)$ at z' is:

$$\frac{\partial Y(z')}{\partial f(z)} = \mathsf{T}(z, z')^{\gamma} \left(Y_t(z')/f_t(z)\right)^{1-\gamma}.$$

The return to f at location z is therefore:

$$\begin{aligned} W_t(z) &= \int_{\mathsf{G}} \frac{\partial Y(z')}{\partial f(z)}\, dz' \\ &= \int_{\mathsf{G}} \mathsf{T}(z, z')^{\gamma} \left(Y_t(z')/f_t(z)\right)^{1-\gamma}\, dz'. \end{aligned} \tag{4}$$

From (2), Y at each z is homogeneous degree 1 in f. Thus, compensation (4) multiplied by $f(z)$ and integrated across G exactly exhausts total global production.

Equation (4) provides, in a partial equilibrium sense, some intuition for the location decisions below. Holding $\mathsf{T}(z, z')$ and $f(z)$ constant, the higher is $Y(z')$, the larger the immediate reward at z. Since T is maximized when $z = z'$, this is therefore an incentive for locating where economic activity is most intense.[3] The denominator, on the other hand, makes it unattractive to be where a lot of f already is.

Two considerations influence a producer's dynamic location decision. First is the adjustment cost (3). Second is the comparison through time of returns locally $W_t(z)$ and returns elsewhere $W_t(z')$. A convenient way to model these influences is to define the benchmark return

$$\overline{W}_t \stackrel{\text{def}}{=} (2\pi)^{-1} \int_{z' \in \mathsf{G}} W_t(z')\, dz', \tag{5}$$

i.e., the *average* return at time t, and then assume that a producer at t in z solves

$$\forall\, t \geq 0: \quad \sup_{\{f_s(z): s \geq t\}} \int_{s \geq t} e^{-\rho s} \left[\left(W_s(z) - \overline{W}_s \right) f_s(z) - C(\dot{f}_s(z)) \right] ds \tag{6}$$

subject to (3), (5), and the given initial density f_t. In (6), producers maximize the present discounted value (discounted at rate ρ) of excess returns less costs of adjustment. They do this by choosing a time path $\{f_s(z) \geq 0 : s \geq t\}$.

That the benchmark return (5) is an unweighted average across G is not critical for the results. What matters is that, rather than say

[3] If T displayed cycles, perhaps with equal peaks eight timezones apart, corresponding to patterns of human sleep, then locating that many timezones apart would be desirable. However, the spatial cycles below while consistent with this behavior, will not in general arise because of this kind of effect.

just the maximum, the *entire* spatial profile of W enters the benchmark return. This together with the radial symmetry throughout the rest of the model will figure importantly in the equilibrium.

Following the usual "stable roots backwards, unstable roots forwards" solution procedure (Sargent, 1987), the optimal decision rule for producers satisfies

$$\dot{f}_t(z) = \zeta^{-1} \int_0^\infty e^{-s\rho} [W_{t+s}(z) - \overline{W}_{t+s}] \, ds. \tag{7}$$

The producer at z increases $f(z)$ if current-location returns are, in a forward-looking present discounted value sense, better than those elsewhere. This response is stronger, the larger is that location premium. On the other hand, the higher is the adjustment cost parameter ζ, the smaller in absolute value is $\dot{f}_t(z)$ in (7), and thus the more sluggish the response.

2.3 Equilibrium distribution dynamics

The economy begins at time 0 and proceeds forever. An *equilibrium* is a mapping X from $[0, \infty]$ to the space of probability densities on G, such that when $f_t = X(t)$, using (2), (4), and (5) in (7) recovers X. In pictures, an equilibrium describes a spatial density in factor inputs for each time instant between 0 and infinity. A *steady-state equilibrium* is an equilibrium that is constant, i.e., $X(t+s) = X(t)$ for all $s \geq t$.

The uniform density $f_t(z) = (2\pi)^{-1}$ is always a steady-state equilibrium: It implies $Y_t(z) = (2\pi)^{-1} \|\mathsf{T}\|_\gamma$ and thus $W_t(z) = \|\mathsf{T}\|_\gamma$ independent of t and z. Using these in (5) and (7) establishes that the uniform density is time-invariant under the model.

The general case, however, is considerably more difficult. The remainder of this paper considers only *Markov* equilibria. By this, I mean equilibria X where there exist operators $T_{t,X(t)}$ indexed by t and X mapping the space of densities on G to the space of their time derivatives such that:

$$\dot{X}(t) = T_{t,X(t)} X(t) \tag{8}$$

and with operator adjoints representable by a stochastic kernel so that for $f_t = X(t)$, we have:[4]

$$\dot{f}_t(e^{i\omega}) = \int_{-\pi}^{\pi} \left[M_{t,f_t}(e^{i\omega'}, e^{i\omega})\right] f_t(e^{i\omega'})\, d\omega' \quad \forall\, \omega. \tag{9}$$

Despite its integral form, the process in (8) and (9) need not be "linear", as the operator T and thus the stochastic kernel M depend on both time t and state f_t.

By the Markov assumption, the right side of (7) changes from an integral over time $s \geq t$ holding z fixed, to one over $z' \in \mathsf{G}$ holding t fixed. That is,

$$\begin{aligned}
\dot{f}_t(z) &= \zeta^{-1} \int_0^{\infty} e^{-s\rho}[W_{t+s}(z) - \overline{W}_{t+s}]\, ds \\
&= \int_{z' \in \mathsf{G}} M_{t,f_t}(z', z) f_t(z') dz' \quad \text{(in Markov equilibrium)} \quad (10) \\
\implies \dot{f}_t &= T_{t,f_t} f_t, \quad (11)
\end{aligned}$$

where I have used the same M and T symbols to highlight that equations (10) and (11) are just the model's equilibrium counterparts of (9) and (8), respectively.

Because of radial symmetry throughout the model, the operator T_{t,f_t} has a *Toeplitz* property, i.e., if it were a matrix, each row would be a circular translation of the one before.[5] But any Toeplitz matrix has for its eigenvalues the discrete Fourier transform of its first row (which changes depending on the matrix) and for its eigenvectors the (unchanging) orthonormal set of complex exponentials

[4] Such a stochastic kernel construction is given in Quah (1997). Futia (1982) and Stokey and Lucas (1989) provide rigorous descriptions.

[5] Grenander and Szegö (1958) is a key reference for Toeplitz operators. Sargent (1987) and Titchmarsh (1962) provide useful expositions of the Fourier tools here. Although different in motivation and mechanism, Krugman and Venables (1997) and Turing (1952) have also exploited these same properties of Toeplitz operators.

$(2\pi)^{-1/2}e^{i\omega j}$, with ω evaluated on the collection of discrete Fourier frequencies $k2\pi/N$, both k and N integer. Thus, when a Toeplitz operator characterizes a dynamic system, the convergence properties of that system can be described using discrete Fourier transforms, and the system's invariants are linear combinations of complex exponentials, or less obscurely, ordinary sine waves.

Of course, neither (10) nor (11) is discrete. Nevertheless, the intuition of the previous paragraph carries.

Under the Markov assumption (8), equations (4) and (5) imply that the expression on the right of (7) depends only the state f_t. Indeed, provided f is a Markov equilibrium, equation (7) gives exactly the stochastic kernel hypothesized in (9).

That all of f across $z' \in \mathtt{G}$ enters in $\overline{W}$, and therefore in (7) has an important implication for the distribution dynamics of the model. To see this, we need first some technical background. One possible attack on the analysis uses partial-differential equation (PDE) methods: Equation (9) implies inter-relations across both time and space—it holds dynamically not just for a single fixed location ω, but simultaneously across a continuum of locations. Analytical solutions are generally unavailable for PDE analyses. The alternative—conceptually easier and that used here—is to view the infinite system of equations (9) as just a single ordinary differential equation (ODE), taking not real values but values in an infinite-dimensional state space of probability densities.

Linearize (10) about the uniform steady-state $\overline{f}(z) = (2\pi)^{-1}$ to obtain:

$$\dot{f}(z) = \int_{\mathtt{G}} \left[\theta_M(z,z')(f(z') - \overline{f}(z'))\right] dz' \\ - \lambda_M \times (f(z) - \overline{f}(z)), \quad (12)$$

with Frechet derivative $\theta_M - \lambda_M I$, the coefficient λ_M a real number and θ_M Toeplitz, i.e.,

$$\theta_M(e^{i\omega'}, e^{i\omega}) = \theta_M(e^{i(\omega'+\omega'')}, e^{i(\omega+\omega'')}) \mod 2\pi \\ \forall M \text{ and } \omega, \omega', \omega''.$$

Consider equation (12). For all θ_M, the eigenfunctions are the complete orthonormal set of complex exponentials,

$$\left\{(2\pi)^{-1/2}e^{i\omega j} : \omega \in (-\pi, \pi],\ j = -\infty, \ldots, +\infty\right\}$$

independent of f, while the spectrum is discrete and comprises the Fourier transform of any one of the sections $\theta_M(z, \cdot)$ (varying with f but independent of z). The proof builds on the calculation:

$$\begin{aligned}\int_{-\pi}^{+\pi} \theta_M(z', e^{i\omega})e^{i\omega j}\,d\omega &= \int_{-\pi}^{+\pi} \theta_M(1, e^{(\omega-\omega')i})e^{i\omega j}\,d\omega \\ &= e^{i\omega' j}\int_{-\pi}^{+\pi} \theta_M(1, e^{i\omega})e^{i\omega j}\,d\omega,\end{aligned}$$

for then $e^{i\omega j}$ is an eigenfunction, and the corresponding eigenvalue is the Fourier coefficient $\int_{-\pi}^{+\pi} \theta_M(1, e^{i\omega})e^{i\omega j}\,d\omega$ (integer j).

To complete the analysis, it is easiest to borrow economists' intuition from a well-known dynamic analysis.[6] Compare (12) with the Cass-Koopmans growth model. For the latter, we know that in the standard case one of the eigenvalues of the 2×2 transition matrix is stable, the other unstable. Initial conditions—configurations of consumption and capital—must be chosen to nullify the unstable eigenvalue for optimality and convergence to steady state. Configurations that do this rely on using the eigenvectors of the transition matrix.

Apply the same reasoning here. Provided θ_M's spectrum—the Fourier transforms $\int_{-\pi}^{+\pi} \theta_M(1, e^{i\omega})e^{i\omega j}\,d\omega$—minus λ_M has negative real part somewhere, the dynamic system (12) has a stable convergent subspace in the space of densities on $\mathtt{G}$ (see Fig. 2). Initial conditions f_t located within this distinguished subspace, when run following (12), converge back to steady state $\overline{f}$. Moreover, any element in that subspace can be represented as a linear combination of the (complex exponential, sine wave) eigenfunctions. This doesn't guarantee clusters, however, as even the flat, uniform density can be so written.

[6] The development becomes necessarily sketchy here, to conserve space. Quah (1999) provides complete details.

The question is whether the eigenfunctions that are activated comprise the entire, spanning set of complex exponentials, or are instead a nontrivial, incomplete subset.

The real-valued coefficient λ_M can be shown to be positive and increasing in the adjustment costs coefficient ζ. Thus, if ζ is too large given the elasticity parameter γ, then the convergent subspace could comprise all densities on $\mathtt{G}$. Indeed, for ζ sufficiently large, the entire space of densities might be steady states. If, on the other hand, ζ is too small, then λ_M might never be large enough to produce a nonnull convergent subspace, and equilibrium need not exist.[7]

The interesting case, when ζ falls in an intermediate range, implies a convergent subspace that is neither null nor the entire space. But then a convergent initial condition f_t is necessarily a nondegenerate linear combination of complex exponentials, i.e., displays distinct periodicities or clustering as in Fig. 3.

2.4 Waves in space and time

It helps to be clear what Fig. 3 says and what it does not say. Each waveform spatial density, on which clusters appear, is a single point in the saddlepoint-stable convergent subspace of the space of all densities on $\mathtt{G}$. Each is, thus, a snapshot at a distinct timepoint in dynamic equilibrium. Economic intuition for these, therefore, is not the same as intuition for a static tradeoff between the forces for agglomeration and those for dispersion. Instead, the intuition is that of saddlepoint stability, as in the Cass-Koopmans growth model.

Upon a disturbance hitting the system—as in Cass-Koopmans analysis—the spatial density has to "twist" into a form such that, given particular initial conditions, waveforms of the kind in Fig. 3 appear. Since the underlying geography is homogeneous, it is the nature of the disturbance that determines where peaks and troughs appear. Of course, with no underlying heterogeneity, that will always be true. The content in the analysis here is that *only* certain wave-

[7] Krugman and Venables (1997) sought dynamics where the system diverged rather than, as here, converged back to steady state.

form spatial densities are consistent with convergence back to the uniform steady-state equilibrium density. Along the transition path, moreover, every point, i.e., every snapshot of the spatial distribution, shows cycles.

These cycles in space arise from the eigenfunctions of the Toeplitz operator T_{t,f_t}. There might also be cycles in time. Instead of monotone convergence back to steady state, there could be cycling (in the space of densities) around the uniform density. Whether those arise, however, depends not on the eigenfunctions but on the spectrum, i.e., the eigenvalues of T_{t,f_t}. While both eigenfunctions and spectrum derive, ultimately, from the economic parameters of the model, there could be cycles in space without corresponding cycles in time. But cycles in space always occur, along stable convergent paths back towards steady state.

To summarize, the uniform "flat" equilibrium is the steady-state to which all (stable) cyclical waveform densities converge. If the system is continually perturbed by disturbances, then almost all the time we observe *only* clusters in economic activity. The clusters will typically occur on different points in space through time. They will wax and wane as history unfolds.

3 Conclusion

This paper has developed a model of Internet economic geography, i.e., where transportation costs and spatial separation don't matter.

The key insight (or, really hypothesis) that the model exploits is that even when physical distance is irrelevant, timeliness or timezone connections might be important. In the model, a tension between technology spillovers across time and "sticky" factor inputs produces clustering or periodicities in economic activity across the Equator, but with location along longitudes left undetermined. The same underlying economic parameters determine waveforms in economic activity simultaneously across both space and time.

One of the attractions of a model like that here is the seemingly spontaneous emergence of heterogeneities (or clusters, in this application) when none previously existed in the underlying attributes. In the global model of this paper, initial uniformity appears as radial homogeneity across space. But then, as shown in the paper, dynamic evolution is determined by a Toeplitz operator. Since Toeplitz operators all have spectra and eigenfunctions that can be related to Fourier transforms, an explicit dynamic analysis turns out to be tractable.

Indeterminacy along longitudes together with zero transportation costs distinguish the current analysis from, say, more-conventional economic geography.[8] It is these features that empirical analysis to distinguish the two strands of work would need to exploit. Part of that development could, in parallel with the analysis above, exploit tools of higher-dimensional Fourier analysis to characterize equilibrium distribution dynamics across the surface of a 3-dimensional globe.

References

Arrow, Kenneth J. (1962) “Economic welfare and the allocation of resources for inventions,” In *The Rate and Direction of Inventive Activity,* ed. Richard R. Nelson (Princeton University Press and NBER) pp. 609–625

Fujita, Masahisa, Paul Krugman, and Tomoya Mori (1999) “On the evolution of urban hierarchical systems,” *European Economic Review* 43(1), 209–251, February

Futia, Carl (1982) “Invariant distributions and the limiting behavior of Markovian economic models,” *Econometrica* 50(1), 377–408, January

Grenander, Ulf, and Gabor Szegö (1958) *Toeplitz Forms and Their Applications,* 2nd ed. (New York: Chelsea Publishing Company)

[8] In addition, two further observations are useful: First, the model provides a dynamic rational expectations equilibrium, rather than ad hoc dynamics imposed on top of intricate static equilibria. Second, the analysis does not strongly rely on specific functional forms; the timeliness T parameterization in particular is general.

Krugman, Paul, and Anthony J. Venables (1997) "The seamless world: A spatial model of international specialization," Working Paper, LSE, April

Quah, Danny (1997) "Empirics for growth and distribution: Polarization, stratification, and convergence clubs," *Journal of Economic Growth* 2(1), 27–59, March

— (1999) "Cluster emergence on a global continuum," Working Paper, Economics Department, LSE, London, October

Rosen, Sherwin (1981) "The economics of superstars," *American Economic Review* 71(5), 845–858, December

Sargent, Thomas J. (1987) *Macroeconomic Theory,* Second ed. (New York NY: Academic Press)

Stokey, Nancy L., and Robert E. Lucas (1989) *Recursive Methods in Economic Dynamics* (Cambridge MA: Harvard University Press)

Titchmarsh, E. C. (1962) *Introduction to the Theory of Fourier Integrals,* 3rd ed. (Oxford: Oxford University Press)

Turing, Alan M. (1952) "The chemical basis of morphogenesis," *Philosophical Transactions of the Royal Society of London Series B* 237, 37–72

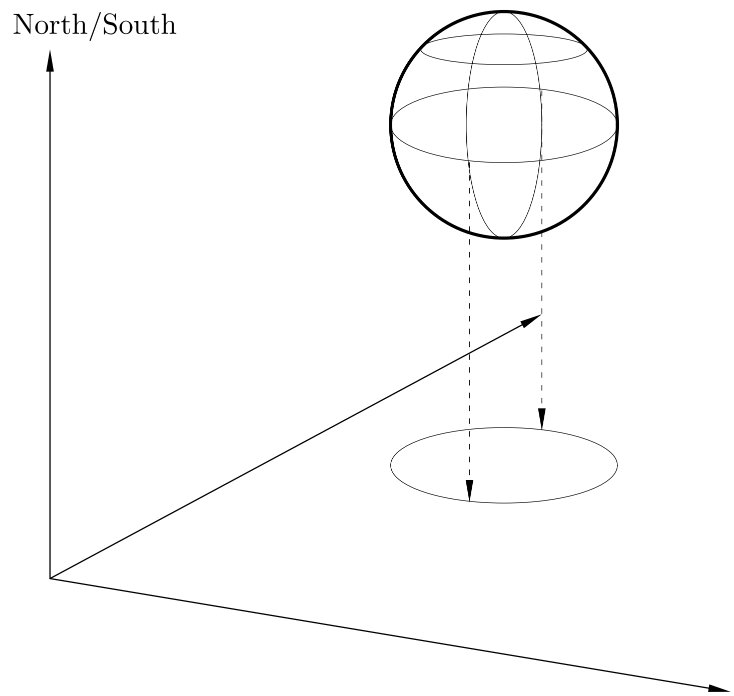

Fig. 1: Collapsing the 3d globe to a minimal isomorphic image Ignoring degeneracies at the North and South Poles, when time matters but not geographical distance, the homogeneous globe is isomorphic with a circle

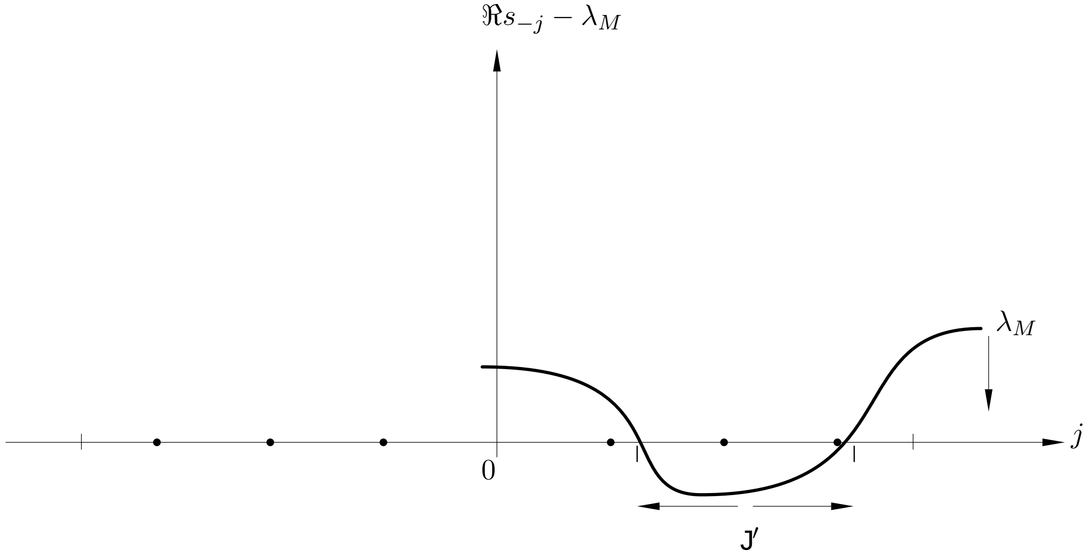

Fig. 2: Real part of the spectrum of θ_M Displaced downwards by increasing λ_M, the positive components that remain activate the associated complex exponentials in the invariant family of eigenfunctions. (The spectrum happens to be discrete, but that is inessential.)

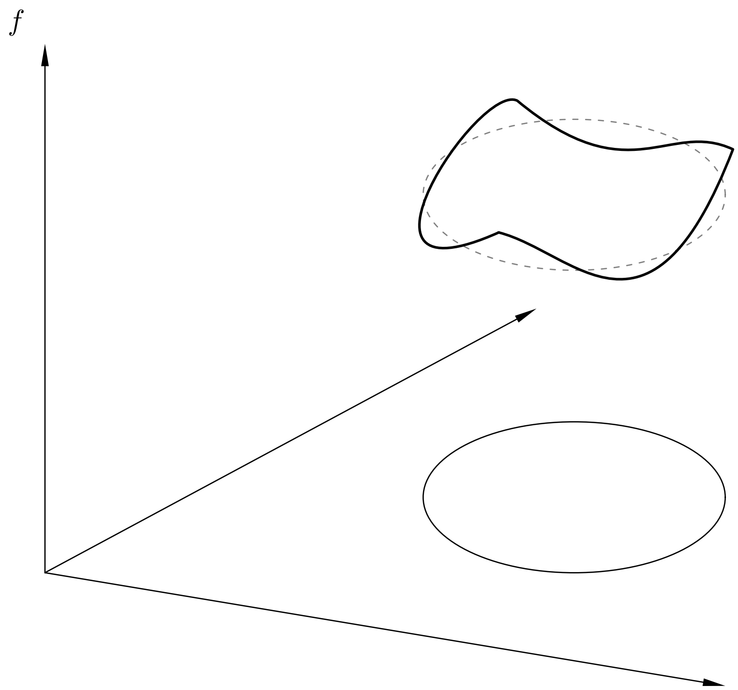

Fig. 3: Cycles in space Local perturbations converge back to steady state only when they display cyclicalities in space