type).

1. Introduction

No matter how finely one disaggregates, there are few markets that sell homogeneous products. Within virtually every market, one can find niches between which there is variation in properties of the goods sold. Automobiles are luxury sedans or sports cars; computers are super or micro, this year's or last year's model; clothes are haute couture or off-the-peg, fashion-forward or 'that timeless classic'; even wheat has different grades. Suppliers, naturally enough, try to meet the demand for differentiated products, producing products that fit the desires of the consumers in the various market segments. From both demanders' and suppliers' point of view, one way that market segments can often be differentiated is through information. What is often referred to as the "high end" of the market typically embodies the latest information—information generated recently to create the latest design, or the latest technology. Goods produced to serve the low end of the market by contrast, typically embody older information—often information that was embodied in yesterday's high end products. This suggests that an important difference between firms serving high and low ends of the market is the activity of knowledge generation.

In this paper we are concerned with the spatial distribution of firm activity. We discuss it in the context of the segmentation of demand along the high end/low end spectrum. At first glance, there seems to be no particular reason why there should be any relationship at all, barring some non-uniform spatial organization on the demand side. We will argue that a relationship is likely to exist, and that it works through features of knowledge generation.[1]

The existence of clusters of production activity has long been recognized—Silicon Valley, Detroit, Chicago have all been used as examples. Recent work, however, indicates that clusters of innovative activity exist independently of clustering in production. Jaffe (1989) finds that the productivity of corporate R&D, as measured by patenting activity, is strongly affected by the local presence of university research. Similarly, Jaffe et al. (1993) find that there is strong localization of patent citation, even after correcting for the localization of production activities) particularly early in the life of a patent. Consistent with the results of Glaeser et al. (1992), they find that geographic space is more important than technology space. That is, while knowledge spillovers appear to be concentrated geographically, they seem to take place both within and across industries. Prevezer and

[1] We take as a working assumption that participation in the high end necessarily involves knowledge or information generating activities. In industries in which this is not the case, then the discussion, presented in terms of the high end/low end distinction, should be re-interpreted as being about knowledge intensive/non-intensive activities.

Swann (forthcoming) also find evidence that agglomeration economies are not restricted to within-industry effects. Indeed, their results indicate that many of the spillovers generating clustering operate between sectors of an industry. They make clear, however, that the details of the nature of the spillovers differ from industry to industry, with the important spillovers in computing existing between sectors within computing, whereas in biotechnology significant spillovers exist between firms and the local science base. Audretsch and Feldman (1994) find that the degree of concentration of innovation in an industry is positively related to the degree to which the industry is information intensive, and that this effect is more important than the degree of concentration of production. The explanatory factor in each of these analyses is the existence of localized knowledge spillovers. Prevezer (1994), however, in a study of entry into the biotechnology industry, finds the existence of both positive and negative externalities. Entrants are attracted to existing research bases, and certain groups of industries within the sector show, as a group, positive agglomeration economies. But it is also the case that many of the industries individually have (weak) negative agglomeration externalities. The presence of both positive and negative forces must be allowed for.

Because of the importance of current information in high end products, research and development will be one of the prominent features of a firm that produces for that segment of the market—an active R&D program is necessary for participation.[2] We would expect, then, to observe in general the types of clustering seen in the empirical analyses discussed above. But while benefits arising from participation in R&D will depend on the presence and strength of knowledge spillovers, they will depend on several other factors as well. In particular, the price of the output, which will be determined in the usual way by considerations of supply and demand in the market for the good, will be crucial. Thus the nature of demand and the nature of competition among suppliers will be important. This implies the possible presence of pecuniary externalities among suppliers, operating through the price of output.

The cost of R&D works in a more complex way and will be affected by both positive and negative externalities. Knowledge spillovers among firms have been widely discussed as a Marshallian external economy. A large pool of knowledge and expertise, even if not being employed directly, greatly facilitates the generation of new knowledge and products. Further, particularly in industries that operate on the knowledge frontier, much of the

[2] We do not hesitate to include design activities as part of R&D.

knowledge employed is tacit, embedded in facilities, structures and people, which means that transmission of it, if it takes place at all, takes place most effectively through face-to-face communication. It is also the case that specialized business inputs can be vital in an R&D intensive industry. Feldman (1994) discusses the importance of patent lawyers, firms doing market research and feasibility studies, testing laboratories, and the presence of appropriate financial capital, the presence of all of which make product development easier. In any industry that has several layers of production, the presence of firms in other layers can be vital both through learning by using activities (Von Hippel, 1988), and through the availability of specialized inputs—the nearby presence of cutting edge button manufacturers is necessary for the success of a cutting edge clothing design house. All of these considerations suggest the presence of local positive externalities among firms in the high end of a particular market.

One source of the negative externalities noted above and observed by Prevezer (1994), may be congestion effects. When Marshallian externalities operate through physical, or through some economic infrastructures, congestion is an obvious possibility. While it is less likely to be highly important for knowledge-based externalities, the potential for disagglomeration economies, often having to do with the competition for specialized inputs in the short run, seems real. If the disagglomeration economies operate at the same length scale as the agglomeration economies, then in understanding the market one need only pay attention to the net effects—are the agglomeration or disagglomeration economies stronger in the appropriate range. If the net effect is attractive, agglomeration will take place (generally speaking), but if the net effect is repulsive, agglomeration of activity will generally not take place. Disagglomeration economies become empirically interesting in their own right, though, if they operate at a different length scale than do the agglomeration economies. If, for example, the agglomeration effects are relatively short-range, whereas the disagglomeration effects operate at longer range, then the second type of effect must be considered independently. Consider as an example tacit knowledge. Tacit knowledge is embedded in facilities, structures, routines and people. Facilities, structures and routines are relatively immobile, and thus spillovers from them will act only over short ranges.[3] Further, they lend themselves to technical rather than pecuniary externalities. People are

[3] Of eight ways of acquiring knowledge, technical workers ranked 'visit knowledgeable person more than 20 miles away' last in preference. (Sweeney, 1987, p. 138, quoting Rosenberg 1967.) The probability of two people communicating at least once a week is about 0.98 percent if their locations are 2 metres apart, but falls to 0.06 percent at a distance of 50 metres. (Sweeney, 1987, p.141.)

another matter, though, in that there is a market for skilled labour. Thus even if there are technical externalities associated with an agglomeration of skilled labour, there will also be, potentially at least, negative pecuniary externalities arising from an agglomeration of firms bidding for that labour as firms attempt to increase their stock of tacit knowledge by hiring skilled labour. Because labour is relatively mobile—certainly beyond the range of "normal" passage of tacit knowledge, and in relatively inelastic supply in the short run, this pecuniary externality will have a longer range than the technical externalities.[4] Summing the externalities that arise from different sources will then produce a net effect that is locally positive but negative at a slightly longer distance. This same argument will apply to any specialized input which contributes to externalities over a short range but is mobile beyond that range. There is the possibility, then, of a regional disagglomeration effect.

There is another way of thinking about regional externalities. In models with non-migrating agents, agglomerations in activity are formed by agents located close to one another doing the same thing. Thus on a more liberal interpretation of such a model the size of an agglomeration or cluster of activity is isomorphic with the amount of activity in a location.[5] In this case, if we are considering production, it is appropriate to think of congestion economies as mediated through distance. From the point of view of an agent in an agglomeration, agents added to the edges of it increase the level of production of that location to such a point that, for example, there is competition for specialized inputs;[6] congestion of physical infrastructure, such as roads, airports, communication systems; or an increase in the cost of living which will raise wages.[7] Thus a regional effect can operate simply as a tailing off, and eventually becoming negative, of the externalities associated with agglomeration.

[4] Of course there may be positive externalities associated with labour mobility, perhaps knowledge that workers carry with them. This would be a positive regional effect. We assume that even if this exists, the *net* regional effect is negative.

[5] This interpretation would be consistent with the results in Prevezer and Swann (forthcoming) suggesting that as employment in an industry at a site increases, entry of new firms at that site slows down.

[6] There were for many years furniture companies dotted all over southern Ontario because of the need for a convenient supply of lumber. But any given stand of trees will only support so many furniture factories.

[7] The high cost of housing in particular is one of the reasons firms give for locating outside Silicon Valley.

This discussion suggests that three types of interactions among firms engaged in this part of the market will be important. There is a global effect through competition in the market for the output. There is a regional effect through the competition for relatively mobile factors of production. Finally, there is a local effect through Marshallian economies. It is the interaction of these three effects that will generate patterns of activity in the production of these knowledge intensive goods.

The alternative to an extensive R&D program and production for the high end market is to participate in the low end of the market. Here, the necessary knowledge and information is of a different nature, as the key to success in this type of market tends to be production efficiency. While there probably are agglomeration effects in this part of the market, they are likely to be much weaker, and to operate in different ways. Global pecuniary externalities from competition in the output market may be important of course.

In this paper we build a simple model of firms in industries in which the market segments along high end/low end lines, assuming that part and parcel of the segmentation is knowledge production. We model knowledge production as one of the main inputs into the production process, and as being subject to the types of externalities discussed above. We use this model of the firm to generate results about the spatial pattern of activity at the macro level. The model is necessarily very stylized, and we make several simplifying assumptions in order to preserve some clarity in presentation and to improve our ability to interpret the results. The results generated by this model must necessarily be qualitative, and indicate patterns, both spatial and temporal, rather then magnitudes, both of steady states and of responses to change.

2. The Model

Consider an economy with N heterogeneous firms, indexed by n, located at fixed points in a two dimensional space. The market in which the firms participate is segmented into two parts—the high end and the low end, but there is no spatial organization of the demand side of the market. At the beginning of every period each firm decides which of those segments to service. We use Z_{high} and Z_{low} for the number (or proportion since the total number of firms is fixed) of firms engaged in the high and low end of the market respectively. The choice criterion is one-period profits. If the firm enters the high end, it performs R&D and then produces one unit of the high end good. If it enters the low end of the market, no R&D is required; the firm simply produces one unit of output.

For firm n, the cost of entering the high end has three parts. There is a fixed cost of production, F_{high}; there are costs of R&D; and there is a firm-specific idiosyncratic component.

Costs of R&D are affected by two things. There are local externalities, $f(n, z_{high}^{local})$, derived from the activities of firms near to n which are also engaged in the high end. There are also what we call regional externalities, $g(n, z_{high}^{regional})$, derived from firms somewhat further away, also engaged in the high end. The idea here is that any firm interacts in two ways with every other firm, but the strength of the interactions changes with distance. In particular, both local and regional effects diminish with the square of the distance between two firms, but the local effect diminishes faster than does the regional effect. Thus we define

$$f(n, z_{high}^{local}) = \sum_{n'=1}^{N} I_{n'} A \exp\{-d_{n,n'}^2/a^2\},$$

and

$$g(n, z_{regional}^{high}) = \sum_{n'=1}^{N} I_{n'} B \exp\{-d_{n,n'}^2/b^2\},$$

where $d_{n,n'}$ is the distance between firms n and n', $I_{n'}$ is an indicator function taking the value 1 if firm n' is in the high end of the market, and 0 if it is in the low end, and A, a, B and b are all constant parameters with $a \leq b$.[8]

There are two potential sources of firm-specific, idiosyncratic cost considerations in this model. Each firm will, over its history, develop its own set of competencies. These will reside in or be associated with, for example, the particular set of technicians or designers employed; the firm's relationships with other firms, both upstream and downstream; its raw materials expertise, and so on. If we assume that tastes shift randomly in the high end of the market, as would especially be the case, for example, in markets where design or fashion is important, then how well a firm's set of competencies is suited to the current state of demand will change randomly. This can be modelled as an idiosyncratic cost component, beneath the vision of any analyst, and stemming from firm heterogeneity, since different firms will have different histories and thus different competencies.

[8] These definitions of $f(\cdot)$ and $g(\cdot)$ indicate that their arguments actually include the states of all other firms in a weighted sum. For pedagogic reasons we continue to refer to these arguments rather loosely as z_{high}^{local} and $z_{high}^{regional}$, the local and regional proportions of firms engaged in the high end market.

Secondly, firms' decisions are based on expectations of the future with regard to the nature of this year's fashions and the difficulty of producing them; beliefs about the processes by which appropriate R&D will be successfully performed; perceptions of this year's demand (both in terms of quality and quantity); and on beliefs about the number of different ways of serving the high end market (if there are many different ways of serving it, it is likely that a firm will choose a way that is different from his neighbours, so the relative the value of spillovers will be reduced). Unless expectations about these sorts of things are generated in a purely mechanical way, they will introduce heterogeneities among firms that will have an effect on expected costs.

These idiosyncratic features of firms, which give rise to heterogeneity, will in general change from period to period. These two sources of heterogeneity are clearly beneath the vision of the analyst and so are modelled simply as a random variable, h_n^{high}.

Benefits from entering the high end of the market arise from selling the product in a global market. The market demand curve is downward sloping, represented by $P_{high}(q)$. This price will be affected by the quantity supplied, or equivalently by the number of firms engaged in this market segment, Z_{high}, since each firm produces one unit. For simplicity, firms remain myopic in their expectations about price. (Time subscripts are omitted for notational convenience only. Clearly, in equilibrium the number of firms engaged in each market segment will not change, so prices will not change over time.)

Net benefits to firm n from entering the high end of the market, then, are written as

$$\Pi_{high}(n) = P_{high}(Z_{high}) - F_{high} - f(n, z_{high}^{local}) - g(n, z_{high}^{regional}) - h_n^{high}.$$

Entering the low end of the market is less complicated. We assume there is no information generation needed, so firms consider only the fixed cost of production, F_{low}, and a demand curve $P_{low}(Z_{low})$, plus an idiosyncratic effect, h_n^{low}. We can write

$$\Pi_{low}(n) = P_{low}(Z_{low}) - F_{low} - h_n^{low}.$$

Thus the firm's chosen action will depend only on the sign of

$$P_{high}(Z_{high}) - F_{high} - f(n, z_{high}^{local}) - g(n, z_{high}^{regional}) - h_n^{high} - P_{low} + F_{low} + h_n^{low}.$$

There are only two activities, so $Z_{low} = 1 - Z_{high}$. Assume linear demand curves and define $\bar{P}(Z_{high}) = P_{high}(Z_{high}) - P_{low}(1 - Z_{high}) = \alpha - \gamma Z_{high}$, where α and γ are parameters. Define also $\bar{F} = F_{high} - F_{low}$, and $\bar{h} = h_n^{high} - h_n^{low}$.[9] Now the firm's action depends only on the sign of

$$\bar{P}(Z_{high}) - \bar{F} - f(n, z_{high}^{local}) - g(n, z_{high}^{regional}) + \bar{h}_n. \tag{1}$$

There are two standard ways of approaching a problem of this nature. One is to solve the N firm-problems simultaneously. This is clearly going to be non-trivial. The second is to look for equilibrium conditions at the macro level, impose them and then try to recover from them properties of firm behaviour in equilibrium. In a standard analysis this would commence with the observation that in equilibrium the marginal firm must be indifferent between its two options. This will define an equilibrium value for Z_{high}, call it Z^*, such that (1) is equal to zero.

This is illustrated in Figure 1. Here, PP is the demand curve, and C_a and C_b are the cost differentials for firms a and b. Suppose that we observe Z_0. For firm a, the cost differential is greater than the price differential so it chooses the old technology. The reverse is true for firm b. The choice of firm a tends to decrease Z; firm b's choice tends to increase it. At the equilibrium level of Z, the expected motion of Z and P are both zero, but out of equilibrium they will drift, as firms make their choices. But because Z drifts, C drifts due to global externality effects. C also drifts if firms "near" firm n are changing their choices. The two sources of drift indicate the complexity of the dynamics of the model. And as a result, in this model the equilibrium condition is not as nice a condition as it appears. For any value of Z_{high} there are many possible configurations of z_{high}^{local} and $z_{high}^{regional}$. Thus even if $Z_{high} = Z^*$, there may be locations in the space where (1) is non-zero, and thus at that location there is motion in the sense of firms changing their activities in ways that might change Z_{high}. This suggests the possibility of many possible equilibrium values for Z_{high}. It suggests further that they may very well be difficult to find. That having been said, there seems to be one immediate conclusion that can be drawn, which is that the bigger the market for the high end good relative to that for the low end good, measured by α, the more firms will do R&D.

[9] This setup can be generalized in the obvious way to allow for local and regional externality effects for suppliers in the low end of the market, to generate $\bar{f}(\cdot)$ and $\bar{g}(\cdot)$.

3. A general class of models

At this point it may be valuable to point out certain central features of this model. The model is essentially dynamic. There are a large but finite number of agents, and they are heterogeneous. There is agent interaction that is non-anonymous, by which we mean that when two agents interact, either mediated through the market or not, the nature of the interaction is affected by the identities of the agents. In particular, the locations of the two agents matter in that they affect the strength of the externalities, and interactions differ depending whether the two agents are engaged in the same segment of the market or not. Finally, there are spatially dependent interactions—local, regional and global externalities, possibly of different signs. This model then, represents a very complicated dynamic process, in which the state space is relatively large, and for which the equations of motion are potentially very complex. In addition, it may be difficult to get information about the exact state of the system since it is large, changing relatively rapidly, and is affected by unobservable factors. These features suggest that a direct attack on it as a dynamic process may, in general, lead to insurmountable difficulties of tractability.

One approach to this class of models sometimes presents a way to finesse the issue of intractability. This is to treat the model as one of two interacting systems. Consider that the pertinent feature of equilibrium is the distribution of economic activity. We can ask, to begin, whether the distribution will have identifiable structure or not. There are, in models of this sort, forces promoting structure, namely the spatially dependent externalities. The type of structure they promote will depend crucially on the exact nature of the externalities, but their existence will lead agents to attempt to coordinate (either by being like or by being unlike) with their neighbours, and this will introduce structure to the activity. On the other hand, though, there are idiosyncratic features, which are, we assume, not spatially dependent.[10] These forces will be introducing randomness into the pattern of economic activity, as agents try to accommodate them. Clearly, if there is an equilibrium it must exist where the two types of forces are in balance.

With this general approach, it is possible to derive equilibrium conditions for this class of models. (This is done in the appendix.) Solving for the equilibrium involves solving a

[10] If there is a spatial structure to the idiosyncratic features, it can be removed and added to the deterministic part of the model, leaving a structure-free idiosyncratic force.

state equation of the form

$$\Pr(\chi) = \frac{e^{\beta\Pi(\chi)}}{\sum_{\chi'} e^{\beta\Pi(\chi')}}, \tag{2}$$

where the probability of finding the system in state χ, which describes the action of every firm, is proportional to the exponential of $\Pi(\chi)$, the sum of profits earned by all firms if the economy is in that state.

4. The Nature of Space

One common assumption made in spatial models is that agents are located on a homogeneous plain. This is very useful in many cases, and can be instrumental in employing various analytical techniques. Often, though, when considering location issues geographic features do matter. Political borders often matter, as do mountain ranges and sea coasts. More generally, different regions of the world are differently suited to different activities; making wine, for example. This has predictable effects on the kind of clustering of wine-making activity that is observed. This comment can be interpreted as a statement that the world is, in many ways, divided into regions that exist before any agents make choices about which types of activities to pursue. Some regions are well-suited to some activities and some are well-suited to others. Part of the effect of pre-existing regions of this type is that to some extent communication patterns are also defined before agents make decisions. It turns out that assuming an extreme form of this feature of the world makes it possible to solve some of these models analytically, to determine the types of activities that are pursued in different places.

4.1 Heterogeneous Space

Assume that there are pre-existing districts, within which all agents interact freely.[11] There is, however, no interaction across districts. This is, in effect, a very particular distance measure. Suppose there are local and global externalities as discussed above, but no regional externalities. Districts might be the countries in Europe, for example, (in the bad old days) if language, cultural or political barriers dramatically raise the costs of communicating with agents in other countries. In this case we can derive equilibrium conditions from equation (2) above. Equilibria are found by solving a set of equations of the form

$$z_{r,q} = \frac{e^{-\beta(f(z_{r,q})-P(Z_q))}}{\sum_{p=1}^{Q} e^{-\beta(f(z_{r,p})-P(Z_p))}}, \tag{3}$$

[11] This discussion is drawn from work presented in Cowan and Cowan (1994).

where $z_{r,q}$ is the proportion of firms in district r participating in market segment q, ($q \in \{high,\ low\}$), and Z_q is the proportion of firms in the entire economy in segment q. Here $1/\beta$ is a measure of the degree of heterogeneity in the idiosyncratic effects felt by the agents. (In a more general model, the functions $f(\cdot)$ and $P(\cdot)$ can be interpreted simply as the local and global externality effects.)

In this type of space, where there is only direct interaction among agents in the same district and no direct inter-district communication, the equilibria have interesting properties. If local externalities are positive and global externalities are negative, as would be the case if R&D exhibited Marshallian externalities, but the output of firms is sold in a global market, there are three types of equilibria. They depend heavily on the degree of heterogeneity in the idiosyncratic terms. If heterogeneity is high, that is to say, firm specific considerations in decision-making vary considerably from firm to firm, there is no structure to the equilibrium pattern of activity. These idiosyncratic features effectively determine a firm's action, since a firm gets high value from taking advantage of (or perhaps compensating for) its particular situation. As heterogeneity falls relative to the strength of local externalities, though, the equilibrium changes and order appears. Idiosyncratic considerations become less important, and agents begin to attempt to coordinate to capture positive local externalities. Structure appears in the equilibrium as each district standardizes on one type of activity. If global externalities are strong, each district will standardize on a different market segment.[12] If global externalities are weak, however, another equilibrium is possible. In this equilibrium all districts standardize, but activities are not necessarily equally represented at the global level. Indeed, it is possible that all districts standardize on the same activity.

The negative global externality, if strong, prevents all districts from standardizing on the same activity. The relatively strong benefits to an agent from pursuing a globally unpopular activity will prevent this outcome. The system can get trapped, though, if the externality is weak, and the benefits to an agent from pursuing a globally unpopular activity are outweighed by the losses incurred from pursuing a locally unpopular activity. This, however, is the only role that the global externality plays. Surprisingly, it has no effect on the degree of local standardization, neither whether it occurs or not, nor the

[12] This statement assumes there are the same number of activities as districts. Generally this is not the case of course, so the result is that districts standardize on different activities to the extent possible, while in terms of global proportions, activities are approximately equally represented.

degree of it (measured say, as the proportion of agents in a district pursuing the dominant activity) if it does.

The degree of heterogeneity, on the other hand, is critical. The tension between concern for idiosyncratic considerations and the desire for local coordination (that is, taking advantage of positive local externalities) determines the degree of standardization of activity. It is not particularly surprising that as the amount of heterogeneity among agents falls, the degree of standardization increases. It is clear that as agents become more and more alike, the benefits to them from coordinating will become more and more important compared to the benefits from accommodating their idiosyncratic competencies. There is, however, a phase change in this relationship. Above a critical amount of heterogeneity, there is no standardization of activities within districts. Below that level as the degree of heterogeneity falls the amount of coordination rises quickly, so that very soon virtually all agents are doing the same thing. There is only a narrow range of values of heterogeneity in which there is partial standardization within a region. Thus in understanding any sort of clustering or agglomeration of activity, the sources and nature of agent heterogeneity will be very important.

This discussion has been based on an extreme topological situation, in which communication patterns are very limited. We return now to the other extreme and consider a space in which there is no pre-existing geography. Agents are located on a uniform plane, and communication patterns are more general, though the importance of communication decreases (or the cost of it increases) with distance. It is clear from the discussion of the model that it is complex, and is likely to be difficult to solve analytically. We solve the system by using a Langevin approach, in which simple dynamics are Langevin equations, the steady states of which are known to be solutions to systems such as those in this model.[13] The system is solved in this way for a variety of parameter values in order to generate robust qualitative results. The discussion in the next sections is based on those results.

4.2 Homogeneous Space

We assume in this model that there are both positive and negative spatially dependent externalities. Following the discussion above, we assume that the agglomeration economies have a shorter range than do the disagglomeration economies. As is common in models

[13] See Van Kampen (1992, Chapter 9) for a discussion of this solution technique.

containing economies of agglomeration, economic activity forms into clusters. The effect of the negative medium-distance externalities is simply to restrict the growth of cities to a certain size. As a city grows it is forced to draw skilled labour from further and further afield. Eventually, it will discover itself in competition with other cities for the labour between them. This will drive up the costs of that labour, and inhibit further growth. The strength and length-scale, in interaction with the other externalities in the model, will determine the upper limits of city size.

A second general comment worth making about models of this sort is that without some form of negative externality or stickiness all agents will eventually perform the same action. We would observe a standard snow-balling phenomenon.[14] David and Foray (1993) generate equilibria having variety by implicitly introducing costs of switching from one activity to another. This will produce a multi-activity equilibrium provided the idiosyncratic considerations can never by themselves overcome the switching costs. That is, when an agent is pulled by the externalities equally hard towards all possible activities the presence of a cost of switching will dictate that he not change his action. But this decision can be overturned if his idiosyncratic cost element is larger than the switching costs. If there is a positive probability of this happening, only a uniform state is stable.

In the model described in section 2 of the current paper, stable equilibria with variety in behaviour exist because of the presence of the global negative externality. This is introduced through the demand side of the two markets. If too many people engage in the high end, the price of output falls, the price of low-end output rises, and it will pay some firms to switch to the low end. If this is the only external effect, though, in this model the stable state involves all high end production taking place at a single location—agglomeration is complete.[15] Any apparent equilibrium that does not have complete agglomeration will, in time, be disturbed by agents' responding to the idiosyncratic factors in their costs, and will stabilize when there is one large cluster.[16] In the model developed here, this outcome is prevented by the presence of regional diseconomies. Cities cannot

[14] See for example, Arthur (1994) chapter 4, or David's (1992) snow-shovelling model.

[15] It should be pointed out that there may be meta-stable states with several distinct clusters that relax only very slowly to the stable state.

[16] This assumes that there is positive probability that idiosyncratic considerations be large enough to cause an agent to defect from his cluster. (An unbounded support of h will obviously do.) The production removed from the economy by this defection must re-appear in the economy somewhere else (in order to maintain equilibrium in the output market) and the lowest cost location would be in the largest cluster.

grow too large, so more than one must exist. Indeed, the model generates states with many cities of different sizes. This result is due to the type of relationship between agglomeration strength and distance ($f(\cdot) = \sum_{n'} I_{n'} A \exp\{-d^2_{n,n'}/a^2\}$). If $f(\cdot)$ were constant to some distance and zero thereafter (and similarly if $g(\cdot)$ were zero to some distance and then constant), all cities would be the same size.

4.2.1 Phase Changes

As in the model of restricted communication discussed above, this model seems to display changes of phase with regard to the degree to which economic activity is clustered, as two of the parameters change value. For large amounts of agent heterogeneity, the primary concern of a firm will be to enter that segment of the market that best suits its particular circumstances. Considerations of coordination with other nearby firms are overwhelmed by idiosyncratic concerns. Thus with high amounts of heterogeneity, there is no structure to the pattern of activity. With low amounts of heterogeneity, though, these two types of considerations change in relative importance. As firms become more and more similar, the advantages to coordinating increase relative to other concerns, so economic activity will be structured, with agents doing similar things located near each other—R&D takes place in enclaves. Between the two extremes is a dramatic change in the amount of clustering. There is a critical value of heterogeneity above which there is no pattern, but below which the presence of clusters rapidly emerges. This is illustrated in Figure 2.[17]

The solid lines in Figure 2 were generated by allowing the system to settle down with a high degree of heterogeneity, and slowly reducing it. (Three such experiments are shown.) At each stage, the system was allowed to settle before reducing the parameter further. The dashed lines were produced by running the opposite experiment. (Again, the results of three experiments are plotted.) The system was stabilized with a low degree of heterogeneity, which was then slowly increased. There is an interesting difference between

[17] Every measure of the degree of clustering has problems but if a phase change occurs using a particular measure of clustering, it will appear in any other measure as well. For the sake of simplicity we have used border length as a measure. The space can be seen as divided into regions (possibly consisting of a single agent) in which agents do the same thing. Border length refers to the total length of the borders between such regions. It is at a maximum if the space is a perfect checker-board, and at a minimum if all agents perform the same action. For a given proportion of high-end producers, it is at a minimum if there is a single cluster with straight-line borders. This measure has the appealing feature of containing a certain amount of economic content in that it is the agents located on a border who suffer losses—they get agglomeration benefits from only some of their neighbours.

the solid and the dashed lines. In the second experiment, there is the apparently counter-intuitive result that as heterogeneity increases border length decreases for a time. That is, as the amount of heterogeneity increases the amount of structure, the degree to which R&D takes place in enclaves, increases as well. This seems odd. The explanation has to do with the ability of a cluster to resist the forces of disintegration. High heterogeneity implies that there is considerable random variation in what agents do. To a small cluster this can be fatal. It is relatively easy for a small cluster to be destroyed by a small number of agents deciding not to produce. This would make costs increase for the remaining agents, and they too stop producing—the cluster disappears. Initially, the activity that was formerly located at that cluster appears at random places in space, but most of this is unstable. The only stable place for activity to appear is within a pre-existing cluster, since this is the low cost location. There have been no changes to the demand side and thus the equilibrium level of production is approximately constant, so if the small clusters are unstable and disappear, the activity must eventually settle in a large cluster. Thus the large clusters grow, and this increases the amount of structure in the economy. Eventually, though, heterogeneity becomes strong enough that even the largest clusters cannot resist, and they too disintegrate.

Because the degree of clustering is determined by the tension between heterogeneity and local positive externalities, a phase change is visible in the latter parameter as well. We can explore the effects of a reduction in communication costs by increasing the distance over which the spatial (local and regional) effects are felt. Increasing a and b (in the functions $f(\cdot)$ and $g(\cdot)$) implies that firms farther away will affect the costs incurred by any firm. "Local" effects extend further. Holding all other parameters constant and increasing the distance over which the local and regional effects operate produces Figure 3. The same sort of effect is visible here as was visible in the heterogeneity parameter. When spatial effects are very localized, (when communication and transportation are expensive) there is no clustering—idiosyncratic effects dominate agglomeration economies. As the range of spatial externalities increase, though, at a critical value structure begins to appear, and the system moves rapidly to a new phase in which firms doing R&D are congregated in a small number of locations.

4.3 Changes in Market Size.

As the world economy becomes more integrated, particularly within the emerging free trade zones, firms' market sizes are changing. Whether the integration will increase or decrease the market size of an industry depends, of course, on the state of that industry relative to the industries in its trading partners. We can use this model to explore the effects both of an increase and a decrease in market size.

4.3.1 An Increase in Market Size

After an increase in market size, production levels increase. It does not happen in a single step, however. Production makes a large initial jump, but then commences a slow increase. The first stage of the change occurs simply because with a larger market, it is profitable for more firms to produce. These firms emerge at many points in the space. Firms that emerge in isolated locations, though, find R&D relatively expensive since they do not have any of the benefits of agglomerations. What happens in the second stage of the production increase is that these producers disappear. Effectively we see a re-agglomeration in which the clusters of activity grow and the isolated producers disappear. This re-agglomeration lowers the costs of the average firm, and thus the market can support more production. This re-agglomeration takes place over a much longer time scale than the initial response to the increase in market size. This is shown in Figure 4.[18]

Figure 4 also shows the response of border length to the same sudden increase in market size. An increase in market size implies that the economy can support more high end producers. If fewer than half the firms are in this segment, then in general increasing their number will imply an increase in border length. But border length overshoots. This is caused by the fact that unstable clusters form initially following the increase in market size—firms spring up all over the place—and then disappear, as they are driven out of the market by the lower cost firms located at larger clusters. Some of the production of these high cost producers will be replaced, and, again, the low cost location is in a large cluster. Adding a firm to an already existing cluster in general has a much smaller effect on border size, so border size falls again to its new (higher than original) level.

[18] Parts a and b of figures 4, 5 and 6 are time series plots of two variables, total production and a measure of clustering, following a shock to the economy. Part c of the figures is a series of depictions of the space of agents after the shock. (They can be thought of as stills from a movie.) On a 50 $\times$ 50 (for a total of 2500 sites) toroidal space, there are nine firms at each site. In the figure, each site is one of 10 shades of grey, indicating how many firms at that site are engaged in R&D and high-end production. Black indicates that all firms do R&D, white indicates that none do. The number of periods (iterations) that have elapsed since the shock is shown beneath each figure.

4.3.2 A Decrease in Market Size

A decrease in market size has a different response pattern. After a fall in the size of the market, as expected, there is a dramatic fall in the amount of production (and in this model the number of producers). Initially, levels of production fall everywhere. This reduces the agglomerations effects, which is fatal to the smallest clusters. Costs of R&D increase at those locations, remaining producers become non-viable, and the cluster will disappear. Now there is too little production, so production at some other location must increase. Again, the low cost location is as part of a large cluster, so the activity that has disappeared from the small clusters, to the extent that it reappears, will do so in a stable way by increasing the size of the larger clusters. Thus re-agglomeration causes production levels to rise after the initial catastrophic fall.

Border length shows a similar over-shooting. The collapse of the small clusters, and the reduction in size of the large ones implies a fall in border length, as shown in Figure 5. A re-agglomeration process increases the total amount of production from this low point, and so increases the number of firms, increasing border length. Interestingly, before, or perhaps as part of, the re-agglomeration process, firms spring up and disappear everywhere in the state space. Figure 5 actually shows a second overshooting. After the first reduction in size, new firms spring up to take up some of the production lost by the collapse of the smaller clusters. In doing so, though, many isolated, high cost producers emerge which will increase border length. They disappear, though, and their production is gradually absorbed by large, stable, low cost clusters.

4.4 An Increase in Attraction Length

Over the course of history communication and transportation costs have continually fallen. Recently, though, there seems to have been a dramatic decline, at least in the cost of communication. To examine this sort of event the model was stabilized at a high communication cost, which in its terms, indicates a small radius of local influence for each firm. This radius was then suddenly increased. The responses of total production and clustering are shown in Figure 6.

A longer reach for local agglomeration economies is equivalent to introducing a way for production costs to decrease. Even if there is no change in the locations of production, some small clusters located near each other will suddenly be receiving some benefits of agglomeration, and their costs will fall. (Some will also have their costs increased by coming into the radius of regional disagglomeration economies, but in general this effect is smaller than the former.) Thus even with no adjustment to the new situation, production can be expected to increase. The possibility of relocation of production, again a re-agglomeration process, will reduce costs to a much greater extent, and production will rise dramatically. This is seen in Figure 6a. Figure 6a suggests a cyclic approach to the new equilibrium production level, but without further exploration of the model it is difficult to tell whether this is qualitatively robust, or merely an artefact.

Generally, in much supply and demand analysis, a decrease in production costs has an effect very similar to an increase in demand. This is not so in the case of the behaviour of clustering, however, as is shown in Figure 6b. In the experiment with increased demand, border length showed an overall increase, as more firms joined existing clusters, making them bigger. When the radius of attraction increases, though, border length decreases. This seems to happen in three stages. First, there is an increase, as very small clusters form, as it is easier to form a cluster, and as firms join existing clusters at random places on the edges. The next stage is that some of these small clusters disappear, as they are, in fact, relatively high-cost production locations. At the same time, existing small clusters join to make larger ones, as the agents between them benefit from both, and are able to reduce costs and become viable producers. This produces a meta-stable state, and one that erodes gradually, as the smallest of these clusters collapse under pressure from lower cost R&D taking place in larger clusters. The production from the collapsing clusters is added to larger, more stable ones, and border length falls to its new stable value.

The general effect of an increase in attraction radius is that the number of R&D enclaves decreases and the average size increases. The number of producers also increases. Neither process is a steady one though. Both happen in fits and starts, possibly including some cyclic behaviour, at least in production. In effect, what happens is that initially firms are added to the edges of existing clusters. (They are added in a random, ragged way, however, which generates an initial increase in the border length measure.) Smaller clusters remain at a disadvantage, though, and they are eventually pushed out of existence because of the larger number of firms now in larger, lower cost locations. But there seem to be threshold effects. Initially a group of small clusters will disappear, reducing border length and production. Production recovers as new firms join larger clusters, and the system stabilizes, temporarily. Eventually, though, there is a new wave of disappearing locations as the (now) smallest clusters are forced out of existence by larger ones. Production falls again but then recovers to a new higher level. These waves of disappearances will continue until the largest clusters begin to reach the ceiling size imposed by decreasing returns.

5. Conclusions

It is well-known that decisions about the location of economic activity are very complex, and in general, involve complicated general equilibrium considerations. Models are necessarily very stylized in order to achieve some level of tractability. The model in this paper is no exception on that score. It is stylized, but nonetheless produces a very rich set of results. In general we find that R&D will take place in enclaves unless firm-specific considerations are very powerful relative to the strength of agglomeration economies. There is, however, a regime change as these two parameter values change relative to each other. We also find that responses to shocks to the parameters induce a two-stage adjustment process. There is typically an initial response, followed by a secondary adjustment through re-agglomeration. This often involves over-shooting. We find further, that there is a subtle relationship between the degree to which R&D is clustered and the total number of firms performing it. This relationship can be seen as the system responds to shocks, and agglomerations form or die, thereby affecting average costs and so levels of production.

The richness of the results generated by this model is obvious from closer examination of the figures. Figures 2 and 3 suggest that there may in fact be three regimes as heterogeneity and agglomeration strength change. Figures 6 indicates that the dynamic patterns in response to changes in communication costs may be very interesting, and somewhat troubling for policy makers. If it is true that the amount of clustering changes fitfully, it may be difficult to predict where or how it will end, and whether steps can be taken to slow down or prevent the decay of smaller centres. Both of these observations are suggestive, but it is not yet known whether they are generally robust, or simply artefacts. The modelling strategy presented in this paper does allow us to address these issues though, and in general is effective in analyzing the complexities of location decisions. It indicates a direction that might further be fruitfully explored in the study of technological change and decisions of location.

References

Arthur, W. B. (1994). *Increasing Returns and Path Dependence in the Economy*, University of Michigan Press.

Audretsch, David B. and M.P. Feldman, (1994). "Knowledge Spillovers and the Geography of Innovation" presented at the 1994 Schumpeter Society Meetings, Muenster.

Cowan, R. and W. Cowan (1994). "Local Externalities and Spatial Equilibria: Technological Standardization and the Preservation of Variety", University of Western Ontario, Department of Economics Research Report 9421.

Cowan, R. and W. Cowan (1995). "Spatially Dependent Interactions: A Statistical Approach to Spatial Equilibria, Technological Standardization and Variety", University of Western Ontario, Department of Economic Research Report 9507.

David, Paul A. (1992). "Path-Dependence in Economic Processes: Implications for Policy Analysis in Dynamical Systems Contexts".

David, Paul A. and Dominique Foray (1992). "Percolation Structures, Markov Random Fields and the Economics of EDI Standards", CEPR Discussion Paper No. 326, Stanford University, and MERIT Working Paper 93-011, Maastricht.

David, Paul A. and Dominique Foray (1993). "Marshallian Externalities and the Emergence of Spatial Stability of Technological Enclaves".

Feldman, M.P. (1994). *The Geography of Innovation*, Kluwer Academic Publishers.

Glaeser, E., H.D. Kallal, J.A. Scheinkman, A. Schleifer, (1992). "Growth in Cities", *Journal of Political Economy*, vol 100(6), pp. 1126-1152.

Jaffe, A.B. (1989). "Real Effects of Academic Research", *American Economic Review*, vol 76, pp. 984-1001.

Jaffe, A.B., Manuel Tratjenberg and Rebecca Henderson, (1993). "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations", *Quarterly Journal of Economics*, vol 108(3) pp. 557-598.

Malecki, E. J. (1981). "Science, Technology, and Regional Economic Development: Review and Prospects", *Research Policy*, vol 10.

Pathria, R.K. (1972). *Statistical Mechanics.* Pergamon Press.

Prevezer, M. (1994). "The Dynamics of Industrial Clustering in Biotechnology" London Business School, presented at the 1994 Schumpeter Society Meetings, Muenster.

Prevezer, M. and P. Swann (forthcoming). "A Comparison of the Dynamics of Industrial Clustering in Computing and Biotechnology", *Research Policy.*

Rosenberg, V. (1967). "Factors Affecting the Preferences of Industrial Personnel and Information Gathering Methods", *Information Storage and Retrieval*, vol 3(3).

Sweeney, G.P. (1987). *Innovation, Entrepreneurship and Regional Development*, Pinter Press.

van Kampen, N.G. (1992). *Stochastic Processes in Physics and Chemistry*, North-Holland.

Von Hippel, E. (1988). *The Sources of Innovation*, Oxford University Press.

Appendix

In this appendix we develop the technique with which we aggregate agents and derive the properties of the global equilibrium. This is the argument from which equations (1) and (2) in the text are derived.[19]

Consider an economy of a large but finite number of agents. Each agent is repeatedly choosing among several activities. The agent pursues this activity for one period, then at the beginning of the next period chooses again. The choice rule is to minimize, over the set of possible actions, one-period costs. Costs from any activity are determined by two types of things. First are deterministic factors. These would be the effects of local and regional externalities, and price effects from global pecuniary externalities. The second type are idiosyncratic to the firm—things such as the interaction between the firm's competencies and the R&D required to produce this year's fashions; or the firm's expectations of the actions of other firms. Because these idiosyncrasies are unobservable to any analyst, we can treat them as a random variable.

In the most general form, what we have is an economy in which total costs are determined by the actions of the agents (described by a vector χ), and by the realization of a random variable or random vector (denoted ϕ). If agents' costs are affected by interactions with other agents, the connection between the random variable, the actions of agents, and costs can be extremely complex. In this appendix we develop a technique for deriving properties of the economy, in particular the distribution of actions, under conditions of arbitrary complexity in this relation.

Consider a random variable that is discrete and finite, with Q possible values; a finite number of agents, N; and a finite number of actions. There is an observable macroeconomic variable (aggregate costs, for example) which is affected by, among other things, the realization of a random variable which is, by assumption, not observable. Thus the macro-variable is a random variable. Imagine running this economy T times, recording each time the equilibrium value of the observable random variable.[20] Because the economy

[19] See Cowan and Cowan (1995) for a more detailed presentation of the material contained here.

[20] "Running the economy" could in principle, be very complex and time-consuming. On the one hand it could involve a single draw of the random variable for each agent, allowing agents to react to their draws, and to the actions of the other agents, and waiting for the economy to reach

is finite and discrete, there are a finite number, M, of possible values of the observable. Label these values $C_1, C_2, \ldots, C_m \ldots, C_M$. Define a vector $A = (a_1 \ldots, a_m, \ldots, a_M)$ such that of the T total trials, a_m trials had the value C_m for the macro variable. If we were to repeat the experiment of T trials, we would generate a new A vector. How likely are we to observe each of these possible vectors? One way to answer that question is to ask how many ways each of them could be generated from T trials. This has a simple answer:

$$P = \frac{T!}{a_1!a_2!\ldots a_M!}.$$

P is thus a frequency distribution over the As. As T increases, P collapses to a delta function.[21] We can thus obtain information about the configuration of the economy by finding the maximum of P. There are two conditions to be imposed, however:

$$\sum_{m=1}^{M} a_m = T,$$

there are exactly T trials distributed among the M costs; and

$$\sum_{m=1}^{M} a_m C_m = T\bar{C}.$$

This constraint states that we are restricting attention to sets of T trials that, as a set, have mean value of the observable of $\bar{C}$. A set of T trials is in effect a sample of the possible economies that could exist under a particular probability distribution of the random variables. Different distributions, of course, will indicate different expected costs for the economy. Because we wish to make general statements about average costs and other macro-economic phenomena, we must ensure that the generalizations we develop regarding these dependent variables are all driven by an economy whose underlying data, including

equilibrium. On the other hand, it could involve repeated realizations of the random variable for each agent, agents responding to their draw and to each other, and the economy evolving until it reaches a dynamic, possibly stochastic, equilibrium. In this second case, the cost recorded would be an average cost, taken over a long time period. Essentially what we are doing here is observing that this economy has some large number of possible realizations. We draw T realizations and record data about each of them. As T becomes large, we have a reliable estimate of the nature of the entire distribution of realizations.

[21] See Pathria, pp.53-61, who describes the limit of this distribution as the number of trials gets large: "In this limit the mean values, the most probably values—in fact any values that occur with a nonvanishing probability—become identical!" T must be large relative to the total number of possible outcomes. With N agents each drawing from M possible values, this implies that T is of the order M^N.

the distribution of random variables, entails a particular average value for its macro variable. Thus we limit attention to sets of T trials having a particular average value, in this case $\bar{C}$. Thus for large T, maximizing the log of P, we have a Lagrangian:

$$\mathcal{L} = \ln P - \lambda(\sum_{m=1}^{M} a_m - T) - \beta(\sum_{m=1}^{M} a_m c_m - T\bar{C}).$$

Because T is assumed to be large, we can use Stirling's formula for factorials: $\ln x! \approx x(\ln x - 1)$. Maximizing with respect to a_m yields, for all M,

$$\frac{\partial \mathcal{L}}{\partial a_m} = -\ln a_m - \lambda - \beta C_m = 0,$$

which immediately gives

$$a_m = \Lambda e^{-\beta C_m}.$$

The value a_m is the count of the number of realizations of the economy in which C_m is the equilibrium value of the observable random variable. Solving for Λ and substituting gives us

$$\Pr(C_m) = \frac{\exp\{-\beta C_m\}}{\sum_{m'=1}^{M} \exp\{-\beta C_{m'}\}}.$$

The probability of observing a state is proportional to the exponential of the macro value of that state.[22] It is, of course, straightforward to do this analysis for a vector of observable variables. This would be desirable in the case of a dependent variable being determined by the values of several independent variables.

It is worth pointing out that using this equation we can solve for $\bar{C}$ as a function of β. Furthermore, this function is monotonically non-increasing.

We now explicitly include the relationship between the observable variable, the random variable and the actions of agents. For example, if the macro variable were determined

[22] There is here an implicit assumption that simply counting states is a good way of measuring the probability of their occurrence. That is, there is an assumption that in the absence of any economic considerations, each state is equally likely to occur. This is in fact much weaker than it seems. If we are trying to model a situation in which it appears, on the face of it, that this assumption is broken, we can re-establish it by re-defining states. In effect, in our model we would clone the more likely states to create a group of similar states (in fact observationally identical states) all having the same probability. This will re-establish the equal probability assumption, and generates a model which is observationally equivalent to the world in which this cloning is not performed.

by the actions of the agents, a vector χ, and the realization of the random variable, ϕ, we could write explicitly,

$$\Pr(\phi, \chi) = \frac{e^{-\beta C(\phi,\chi)}}{\sum_{\phi',\chi'} e^{-\beta C(\phi',\chi')}}.$$

Now it remains to extract marginal distributions. The complete state of the economy, (ϕ, χ), is in practice unobservable. Integrating over ϕ though, will give the probability that the vector of actions takes on a particular value, χ. Formally, the integral can be performed for each agent, leaving a total cost of $\sum_n \bar{c}_n(\chi)$, where $\bar{c}_n(\chi)$ is the average cost of agent n when the vector of actions is χ, where the average is taken over possible values of the random variable. We can therefore write that

$$\Pr(\chi) = \frac{\exp\{-\beta \sum_n \bar{c}_n(\chi)\}}{\sum_{\chi'} \exp\{-\beta \sum_n \bar{c}_n(\chi')\}}.$$

In the text, we discussed a situation in which agents maximize profits, so we can use notation in which the macro variable is aggregate profits, denoted Π. In the model with a small number of pre-defined districts, if $q_n = 1$ and agent n is in district 1, $\bar{\Pi}_n(\chi) = f(z^1_{high} - P(Z_{high}))$, since the mean of h is zero. Now we can extract the probability that an arbitrary agent in district r pursues activity q. Since agents within a district are indistinguishable, this will be equal to the proportion of agents in district r pursues activity q:

$$z_{r,q} = \frac{e^{-\beta(f(z_{r,q})-P(Z_q))}}{\sum_{p=1}^{Q} e^{-\beta(f(z_{r,p})-P(Z_p))}},$$

which is equation 3 in the text.

The outstanding issue is the interpretation of β. Consider a simpler model which has no global externalities—the price of the outputs are fixed, and without loss of generality assume them to be zero. Here, $\bar{\Pi}_n(\chi) = -f(z_{r,q})$. In this case,

$$Z_{r,q} = \frac{e^{\beta f(z_{r,q})}}{\sum_{p=1}^{Q} e^{\beta f(z_{r,p})}}.$$

If β is zero, $Z_{r,q} = 1/Q$. All activities are equally common in all districts. If β is large, $Z_{r,q}$ goes to 0 or 1—every agent performs the same activity, namely the one with the highest average profits (where the average is taken over the distribution of idiosyncratic factors). These two results correspond to results that obtain if there is a large or small degree of agent heterogeneity. If agents' unobservable idiosyncrasies are wildly different from each other, there will be little incentive to coordinate with neighbours, and each agent will 'do his own thing'. Since by assumption this distribution has zero mean, each activity will be equally prevalent. On the other hand, if all agents draw exactly the same factor in every period, all will choose the same action. Thus we can interpret $1/\beta$ as being a measure of the degree of heterogeneity present in the system.

On Clustering in the Location of R&D: Statics and Dynamics

Robin Cowan[†] and William Cowan[‡]

June 1995.

[†] Economics Department, University of Western Ontario, London, Ontario N6A 5C2
[‡] Computer Science Department, University of Waterloo, Waterloo, Ontario N2L 3G1

We thank Stan Metcalfe and the other participants of the seminar on "The Economics of Localized Technological Change" in Turin, June, 1995.