

The Dynamic Urban Model: Transport and Urban Development

John Swanson

Steer Davies Gleave

28-32 Upper Ground, London

United Kingdom

SE1 9PD

Tel: +44 20-7919-8500

Fax: +44 20-7827 9850

Email: j.swanson@sdgworld.net

Abstract

This paper describes a model built to simulate the interaction between transport, land-use, population and economic activity in an urban area. It provides a brief discussion of how the need for such a model arose, and of other modelling methods that have been used to address the problem, touching on the contribution of Forrester's original Urban Dynamics.

It describes how Forrester's model has been adapted and extended to allow a zonal representation of an urban area, coupled with a full representation of road and public transport networks allowing choice of mode and route. A discussion is provided about how hierarchical logit models are used to model choices by travellers on the networks.

Finally brief descriptions are given of four applications of the model, three of them consultancy assignments carried out by Steer Davies Gleave, and the other a hypothetical town that has been used to develop and test the model.

Transport and economic activity in urban areas.

The question of whether transport investment contributes towards economic well-being, or more frequently, towards economic regeneration, in urban areas has long taxed economists and transport analysts. The claim is often made in support of proposed transport schemes in the UK: the investment is needed to boost the local economy, to create jobs, to regenerate the area. In fact while the idea that transport should be related somehow to economic activity seems grounded in common sense, the evidence for the nature of that relationship is far from clear. A UK government advisory committee produced a very substantial report on the topic in 1999¹ drawing on research from around the world, and concluding that, at least in developed economies with existing transport systems, the case was not proven. 'Necessary but not sufficient' seems to be the most that can be said about transport investment; some authors even question whether it is necessary.

¹ "Transport and the Economy", The Standing Advisory Committee on Trunk Road Assessment (SACTRA), Department of Environment, Transport and the Regions, 1999.

The topic has of course been a rich source of research and modelling work over the years. Most writers agree that urban economies are complex ‘systems’, but then do not go on the use system models to try to study them. The analytic framework tends to be grounded in equilibrium economics and techniques, even when it is apparent that the problem is essentially dynamic². The earliest large scale transport models, known as four-stage models and still in use, were equilibrium models, operating in four sequential steps: trip generation (estimating how many trips each household will make); distribution (working out where they will go); mode choice (car, public transport etc); and assignment, in which the trips are assigned to routes through the networks. An iterative procedure is used to cycle round until equilibrium, or something close to it, is reached.

Land-use and transport interaction models (LUTI models) take this even further. They usually comprise a transport model, with some or all of the four stages, and a land use model that attempts to allocate households and economic activity to each zone in the model in response to changes in the network, and then to describe how conditions on the transport network will change in response, and so on. Typically the procedure will be something like: describe the land use patterns in the starting year; run the transport model until it reached equilibrium; revise the land use patterns in the light of the new transport conditions in some subsequent year; run the transport model until it reaches equilibrium for that year; revise the land-use patterns, etc.

These models are usually cumbersome and expensive to set up, requiring large amounts of data. They run slowly, and may not in fact reach equilibrium. Their developers often stress the ever-finer segmentation and micro-detail these models³ provide, but for all this their reputation is not very high.

For system dynamicists the objection to such models is not that just they are technically inefficient, but that they seem to fail to address the central issue: that our major towns and cities are in states of constant flux, the product of many interacting forces acting on short and long time spans with feedbacks also operating on many time scales. Their dynamism arises because of tensions caused by non-equilibrium, and it requires a dynamic model, a system dynamic model, to describe the processes. This is what the Dynamic Urban Model attempts to do.

System Dynamics and Urban Economics

Forrester’s Urban Dynamics is well known in the SD community. His model described the growth and eventual stagnation of an urban area over very long time periods of up to 200 years, but it actually said little about transport. There was an implicit assumption that an adequate degree of accessibility existed in the simulated city, but nothing

² As for example in ‘Transport Investment and Economic Development’ by David Banister and Joseph Berechman, UCL Press, London, 2000, page 38. After explaining numerous ways in which travel behaviour can change dynamically, they then say: “ This discussion implies that, in modelling the effects of transport improvements on the local economy, it is necessary to carry it out within an equilibrium framework, so that the changes in total demand for infrastructure facilities, resulting from changes in travel patterns and rates and in spatial location, are equilibrated with network supply.”

³ For instance, we were given access to a LUTI model built recently that had over 30 household categories and 24 types of employed resident. These models are often strong on cross sectional detail like this, but weaker on structure.

explicitly stated. It was essentially a two-zone model: the city itself, and the outside world.

Nathan Forrester has taken the idea further, by constructing a zonal model with transport costs between the zones represented by a function of the straight-line distances between them. Congestion arose as the numbers of movements increased; while each zone was treated rather like its own example of the original Urban Dynamics city.

The Dynamic Urban Model (the DUM) described here borrows many ideas from Forrester's original. The idea of 'attractiveness' as a location in which to live or do business remains, and many of the mechanisms used can be traced back to that model. However where the new model gains is:

- It allows the modelled area to be divided into zones, each of them stocked with people, houses, businesses, etc;
- It provides a full representation of the transport networks in a manner that is consistent with more traditional transport modelling. This includes software written to convert transport network models built using commonly used transport packages into a format suitable for the DUM;
- It makes use of a hierarchical logit structure to handle choices of travel mode, route, and willingness to travel at all;
- It allows multiple classes of households, people, houses, business sectors etc, each with their own preferences;
- It is set up to read initialising data describing the urban area from an external Excel database.

The DUM has been built to be as general as possible, so that it can readily be adapted to new applications.

A Model of Transport and the Economy

If transport is to impact on the economy of a region or urban area, then it must be via the patterns of accessibility it permits, in that it allows movements of people or goods within economically acceptable times and costs.


Several types of accessibility might be expected to apply:

- For the workforce, there is access to employment, or more strictly to employment that matches skills;
- For employers there is access to a workforce, which affects ease of recruitment; and
- Access to markets and suppliers, both within the modelled area and beyond it.

In general we might say that the longer or more expensive a trip, the fewer people will be willing to undertake it. This gives rise to the idea of a deterrence function, relating the cost of travel to willingness to undertake the trip. Figure 1 shows a deterrence function calibrated using the DUM in an application in Merseyside, in England. It relates to travel to work, and shows how the proportion of people willing to commute

falls as the commute time increases. (In fact the implementation of this curve is via a hierarchical logit formulation, discussed further below.)

Figure 1: A deterrence curve



Given a deterrence curve and information about the costs and times of travel between each pair of zones, then the workforce willing to access a given zone, i , is given by:

$$\text{Accessible workforce}_i = \sum_j \text{workforce}_j \cdot f(\text{cost of travel between zones } i \text{ and } j)$$

Here $f()$ is the deterrence function that returns the proportion of people willing to travel given the cost involved⁴. This is simply the workforce in each of the other zones multiplied by the proportion willing to travel to it⁵. Similarly the number of jobs accessible from a given location is given as:

$$\text{Accessible jobs}_i = \sum_j \text{jobs}_j \cdot f(\text{cost of travel between zones } i \text{ and } j)$$

Access to markets and supply chains is rather more awkward. For the retail sector, customers are members of the public, so one obvious measure is the number of people living within accessible range of a site, calculated in a similar way as above. More generally however businesses do business with each other, and business customers and suppliers may be located within the modelled area or outside it.

For each location the numbers of accessible businesses can be calculated in a similar way as above, given a deterrence function. Connections with the outside world are often important, but the model, by definition, does not know anything about the world beyond its own geographic boundary, and for the time being all that can be done is to include large zones whose role is to represent the world beyond the region of interest.

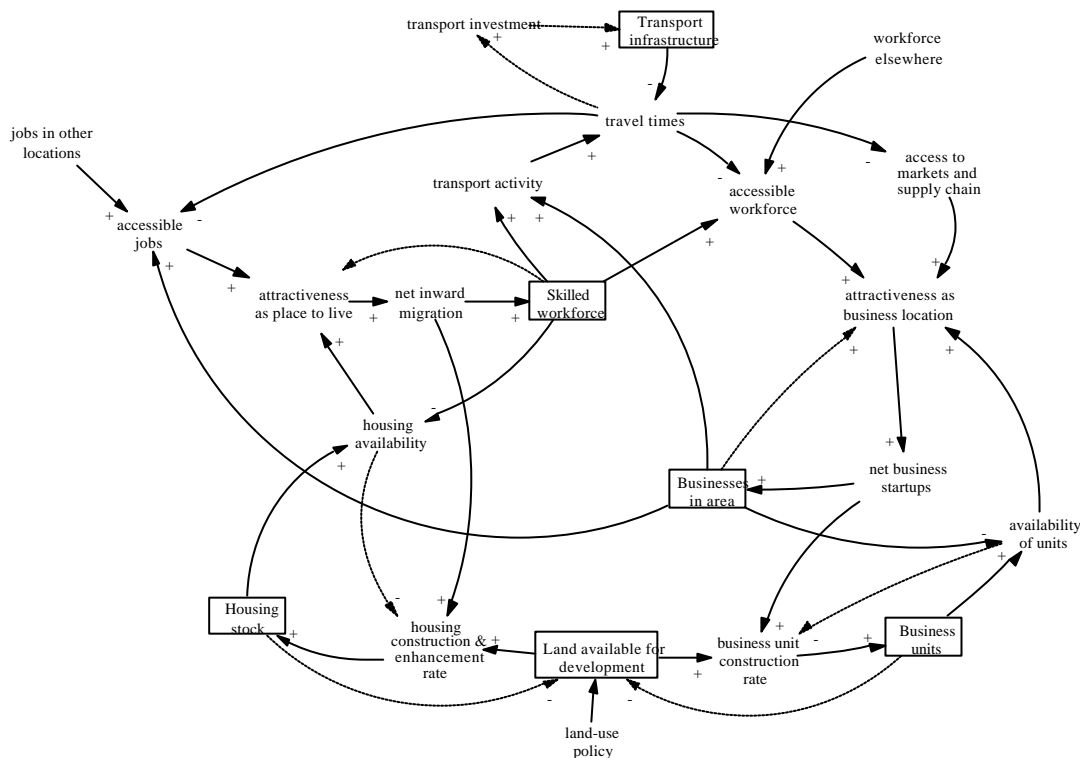
⁴ Some measures of that cost are discussed below

⁵ Travel within a zone is a slightly special case, for here we will not usually have a representation of the transport networks. However it is still possible to calculate accessibility functions within a zone, and the same principal applies.

More information about how transport costs are calculated is given further below, but using this simple notion of accessibility, we can now look at some of the main components of the model.

Figure 2 illustrates some of the primary linkages in the model.

Figure 2: Some linkages in an urban area



Transport is represented here as a stock, right at the top of the picture. If this stock is added to, travel times (more generally, the generalised cost of travel) will fall, so that the number of accessible jobs, from the point of view of residents, will increase. Similarly, access to markets and supply chains will tend to improve. Both of these then affect the attractiveness of the town as somewhere to live and to do business, and consequently the net migration rate and the business start-up rate.

Attractiveness as a place to live is also affected by the availability of housing, and of course if the migration rate increases the availability of housing will fall, unless more is built. Similarly from the point of view of businesses, attractiveness is considered to be affected by access to a workforce, to markets and supply chains, and by the availability of suitable premises.

Finally, at the bottom of the diagram, we have land-use policy and the construction of houses and business premises. The rates at which developers will add to the stock of houses or business units is affected by the availability of land on which to build, vacancies, and their view of expected demand.

The Dynamic Urban Model

The Dynamic Urban Model is built in Vensim, with a supporting database in Excel. The area being modelled is divided into zones, and each zone is stocked with:

- The total land area, the land initially occupied by business premises and by housing, and optionally, the areas of land allocated for business and housing development;
- The numbers of housing units;
- The number of households;
- The number of business premises;
- The number of businesses.

The following sections say more about how each is handled.

The Housing Sector

This is a simplified version of Forrester's model. Housing units are built, age, refurbished and, eventually, demolished. The rates at which these events occur depends upon conditions as they might be seen by developers, and there is a lag between the instantaneous conditions arising and construction activity taking place.

The factors taken into account are:

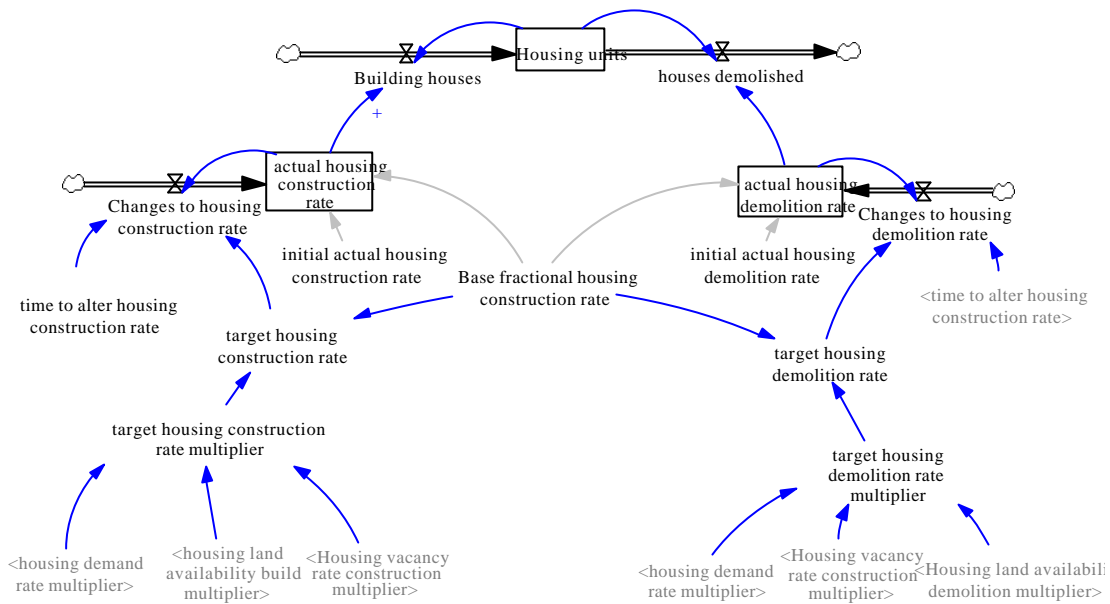
- The available land;
- The current number of vacancies. High vacancies discourage new construction, and vice versa;
- Market trends, and indicated by the smoothed net household migration rate for the zone.

The model can distinguish between different types of housing as necessary. The differences lie in:

- The density at which they are built; and
- The preferences of different types of household for types of house.

A co-flow (not shown) is used to keep track of the average age of the housing stocks in each zone. The interpretation of age is a little loose, for it does not mean the time since originally built, but since last refurbished. If renewal of the stock continues at an adequate rate the average age will remain steady; if the renewal rate falls, the stock starts to deteriorate and become less attractive; if the renewal rate rises the stock improves and becomes more attractive.

Figure 3: Housing construction



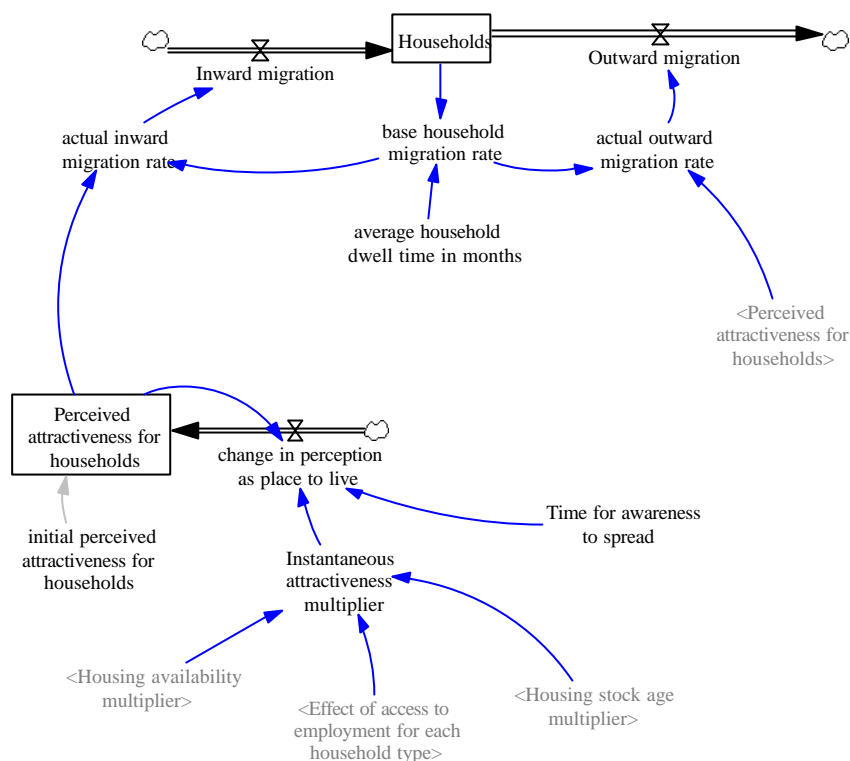
Migration

Again, this is a simplified version of Forrester’s model. Households migrate in and out of each zone, at rates that are affected by how attractive the zone is to live in. Factors taken into account are:

- Access to suitable employment. Although this is derived from the accessibility measures described above, the measure actually used is the average time for people out of work to find employment, which is an output of the recruitment sector described below. The number of accessible jobs or vacancies is not something people are likely to know; the time to find work is something they will experience, or see others experiencing, and they can form at least an approximate estimate of it;
- The availability of housing. The measure is the vacancy rate, which is used as a proxy for price. Each household type has its own preferences for types of housing, and the availability measure is calculated taking this into account;
- The age of the housing stock, generated by the co-flow described above.

A lag is used to introduce delay between instantaneous conditions and people perceiving them and taking action.

Figure 4: Household migration



Different types of household can be defined. They differ in terms of their preferences for types of house, and the mix of people in them. People can be defined in terms of their employment skills, and there is a correspondence between types of people and types of job offered by employers. Not all household members need be in the workforce, that is they are not all necessarily available for work.

Each type of household also has a specified level of car ownership (cars per household) and fraction of its members that have a driving licence. They will also have a specified average monthly expenditure on retail goods.

Construction of business units

Structurally this is similar to the construction of housing. Developers respond to the attractiveness of the zone as a development site, taking into account the availability of land, current vacancies and trends in business activity.

As with housing, there can be several different types of business unit. In practical terms their difference lies in the density at which they are built, the preferences of each type of business for premises, and the average floor space of each.

A co-flow is used to keep track of the average ‘age’ of the stocks of units exactly as with housing.

Businesses

Structurally these are handled in a similar way as household migration. Businesses start up and close down in response to conditions in the modelled area. Different types of business can be defined, each with their own preferences for suitable premises, and profile of employees. The mix of employees is described in the same terms as types of people, above.

The attractiveness of the area as somewhere to do business varies in response to the availability of suitable premises (ie the vacancy rates); the availability of a suitable workforce; the condition of the stock of business units; access to markets and supply chains.

The measure of availability of a suitable workforce is the job vacancy rate. This derives from the accessibility calculations described above, but again it seems better to use a measure employers are likely to have direct experience of, the difficulty of recruiting staff, than a more abstract measure like accessible workforce. Vacancy rates are calculated in the recruitment sector, described below.

Measures of access to markets and suppliers were discussed above. To recap, there are two types:

- For retail, access to residents;
- Access to other businesses located within the model.

In fact for retail the model estimates how the monthly household expenditure will be distributed across the accessible retail floor space, and seeks to expand or contract the retail sectors depending on how the revenues compare to a target required earnings level per square metre per time period.

Recruitment

Figure 5 illustrates the core of the recruitment cycle. This is the mechanism that generates the travel-to-work patterns, and handles the process by which people move between jobs.

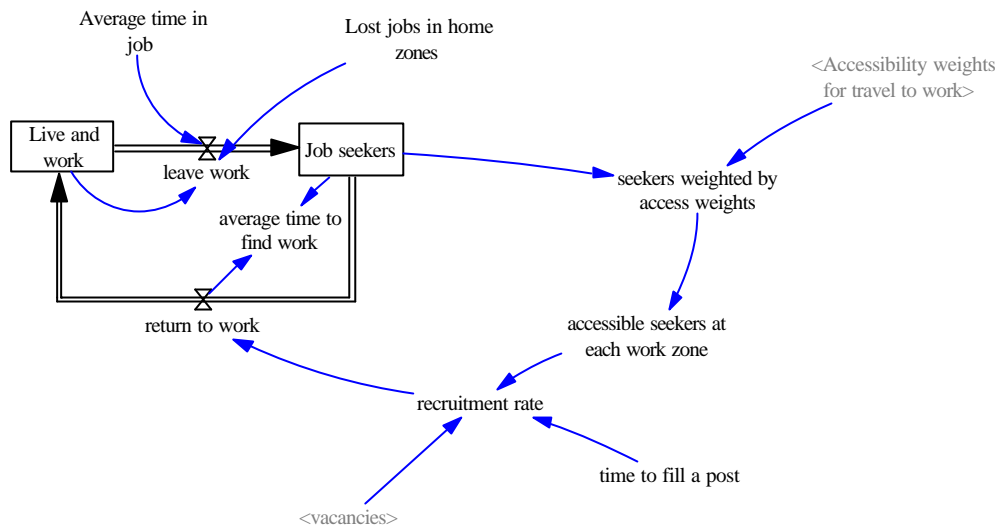
The stock 'live and work' is a three dimensional array, indexed on 'home zone', 'work zone' and person type. It stores information about where people live and work, and is, by implication, the travel to work trip matrix⁶. The 'person type' index allows us to keep track of different types of people, categorised by employment skill.

In time people leave their jobs, either voluntarily or because their employer has closed down. They then join a stock of 'job seekers'⁷.

⁶ Travel to work trips are the only trips explicitly present in the model. Other trips are absent, but accessibility measures are still constructed for them. Retail, and business proximity to other businesses are examples of this.

⁷ The model assumes that everyone leaves their current employment before looking for another job. This is not true, of course, but for the time being is a reasonable approximation. The main implication is that the number of job seekers is an over-estimate of unemployment. The diagram also excludes the effect of migration, which will affect the numbers of people in work and seeking work.

Figure 5: The recruitment cycle



Meanwhile the employers will have vacancies for which they wish to recruit staff. For each employer, there is a pool of accessible job seekers living within range, and the employer will recruit from among them. These accessibility weights are given by the deterrence function, described more fully below. The model recruits job-seekers back into employment from each zone in proportion to their accessibility; it will tend to recruit more heavily from adjacent zones than from those further away.

The effect of this process is that travel to work patterns change as the transport costs change. Reductions in transport costs will tend to increase the range people will travel, and the trip patterns become longer and more dispersed – a pattern that has been seen in urban areas in many countries as car ownership rises and the highway networks are expanded.

This module generates two important indicators:

- The average vacancy rate for employers, which is used as a measure of access to a workforce; and
- The average time to find work, used as a measure of access to employment.

Transport Networks

It is now possible to explain how the model represents the transport system.

In fact the model is capable of representing transport networks in two ways: by zone to zone generalised costs, or by a link based network model.

Most trips, at least by mechanised transport, will involve time and money. Generalised cost is a commonly used measure that combines both of them linearly as in:

Generalised cost = monetary cost + travel time * value of time.

The 'value of time' converts time to money 'equivalent'. It represents the rate at which people trade off money and time, and is widely used in transport analysis. Research methods are available to give estimates of the value of time, which can differ by type of person and type of trip.

Sometimes generalised time is used instead. This is given by:

Generalised time = travel time + monetary cost / value of time.

More generally, other terms might be included. Public transport trips for instance might involve several stages: walking, waiting, travelling, interchanges etc, and these can all be incorporated into the generalised cost or time.

All transport modelling packages can calculate the generalised times or costs between all pairs of zones, differentiating between car and public transport. These can be read directly into the Dynamic Urban Model.

A moment's thought will reveal that this way of working loses sight of individual links in the network, with two immediate consequences. First although we will be able to estimate movements between pairs of zones, we cannot say what volumes of traffic flow on any given link, and, consequently, it is not possible to build in any congestion or crowding response.

Second, we cannot change the characteristics of individual sections of road, or test the imposition of such things as link based tolls within the model. This can only be done by going back to the network model and recalculating and outputting the generalised cost matrices again.

However the Dynamic Urban Model is capable of representing public and private transport networks, link by link. It does not have a path building routine, but rather the available routes between any pair of zones must be defined externally as a sequence of links. This is coded in the model in an array, as in:

Route[i,j,link_no,route_no] = 1 if the route number 'route_no' between zones i and j contains the link 'link_no', and zero if not.

This is fairly tedious to set up by hand, but we have developed software that can take the outputs from conventional transport models, which are very good at path finding through networks of many thousands of links, and write it out in a format the Dynamic Urban Model can read directly. Each link also has a set of properties: its length, free flow speed and capacity for roads, plus such things as frequency, fare and capacity for public transport.

This works well, but of course slows the model down significantly for a network of any size, and arguably introduces a level of detail that is not necessary for longer term

studies. The choice between the two types of model depends on the purpose for which the model is to be used. For longer term studies designed to identify transport strategies that will best aid the local economy, the generalised cost matrix approach is probably adequate; for shorter term studies in which the impacts of specific network changes are to be examined, the network models are needed.

Logit Models

Logit models are commonly used to describe how choices about transport are made, such as choice of route and of mode. Logit models propose that given a set of available alternatives, the probability of any one of them being chosen is given by:

$$P_i = \frac{\exp(U_i)}{\sum_j \exp(U_j)}$$

where U_i is the utility of alternative i . Utility is related to the attributes of each alternative, which in the context of transport will be the characteristics of the journey: time, cost etc. In practice we cannot do more than postulate a functional form for utility, and most commonly a linear form is chosen, such as:


$$U_i = a.Time_i + b.fuel\ cost_i + c.parking\ charges_i$$

where a , b and c are parameters to be estimated⁸. Of course other terms might be included if relevant: quality of road surface, reliability, congestion etc have all been used in transport studies.

The appeal of the logit model is that given a set of alternatives, it returns a probability of each being chosen and that rises or falls as the utility rises or falls. Figure 6 gives an example of a logit choice model, where the choice is between two alternatives, A and B. The horizontal axis is the difference in the utilities of A and B, while the vertical scale gives the probability of choosing A. As the utility of A rises, relative to B, the probability of choosing A rises. Parameters can be chosen to vary the slope of the curve and the point at which it crosses the vertical axis.

⁸ This is, of course, a multiple of the generalised cost.

Figure 6: A logit curve.



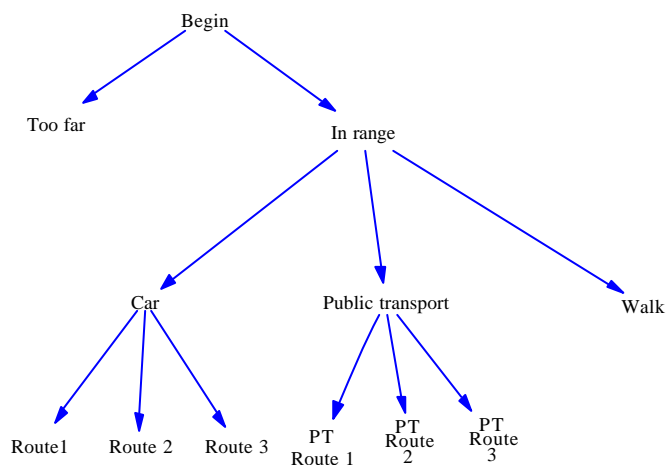
An extension of this idea is the hierarchical logit model. This applies where alternatives can be grouped together in clusters, or ‘nests’ of similar alternatives. Indeed these models are often called nested logit.

To see how this can be used, consider Figure 7. This illustrates a hierarchy of choices that might be considered to apply for someone considering travel between two locations. Right at the top, there is a decision about whether the trip is simply too far or too expensive, taking into account all the available transport options. The model will return a proportion of people who would or would not be willing to make the trip, given the transport conditions.

At the next level, for those who are willing to make the trip, is a choice between mode, in this case between car, public transport or walk. In each case the choice is made taking into account the attributes of each available route.

Finally, at the lowest level, having chosen a mode, there may be a choice of which route to use, out of those available (except, in this case for walk).

Figure 7: The Hierarchical Logit Structure for Travel Choices



Right at the bottom, the choices between routes can be modelled using a logit model as described above. Above this the choices are handled by a further logit model that responds to a measure of the total utility of each mode assessed over all available routes⁹. This measure, known as the logsum, is given by:

$$Logsum_i = \mathbf{q} \cdot \ln \left(\sum_j \exp(U_j) \right)$$

This gives the logsum for branch i , where j counts across all the sub-options available in branch i , and \mathbf{q} is a parameter whose value lies between 0 and 1. It can be regarded as a utility for each available mode. Finally a similar formulation is used right at the top to describe the proportion of people willing to travel or not.

The model uses this structure to model decisions about travel to work. The branch at the top of the logit tree is especially interesting because it is used to generate the deterrence function for travel to work, as shown in Figure 1. The effect is that for any pair of zones, as the travel times and costs increase fewer people are willing to make the trip, and they fall out of recruitment range for employers; conversely, as those times and costs fall, more people come into range.

These models require a number of parameters, but these can be obtained in several ways. First, many of them have been borrowed from earlier research. For example, the ‘value of time’, which is the relative weighting given to travel time and money, is extensively researched, and suitable values are known.

Second, some parameters have been chosen on the grounds of the logit models they give rise to. For instance, it is relatively easy to plot the logit curves in Excel, and select parameter values that give reasonable curves.

Third, several parameters were chosen so that they generated results in the SD model that were consistent with known information, such as travel to work matrices. In fact the Dynamic Urban Model has been used to calibrate deterrence curves for travel to work, essentially by adjusting the parameter values incrementally until travel to work patterns and mode choices are generated that correspond closely to those actually observed in practice in the UK.

A dynamic logit model

An obvious criticism of the logit model as described above is that it supposes that everyone has full information about all the alternatives available to them, that this information can be specified in advance, and that choices are made immediately. In practice none of these things is likely to apply. People do not necessarily seek out information about all the options open to them unless they have to, and in transport the

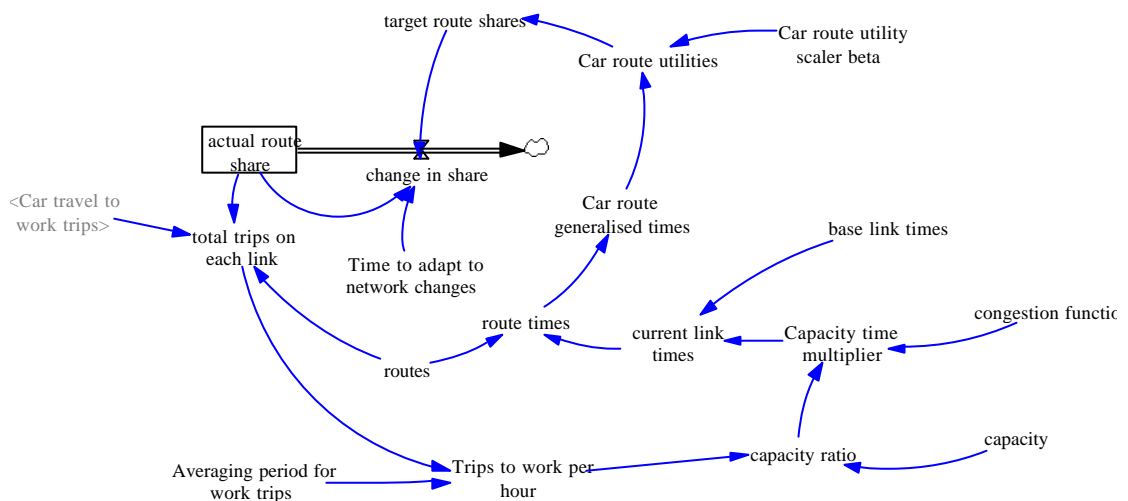
⁹ See for example the Ben Akiva and Lerman reference for a very extensive discussion of these models and the theory underlying them.

choices they make are likely to impact directly on the conditions experienced due to congestion or crowding.

The model handles this by monitoring two sets of choice proportions: the ‘ideal’ proportions predicted by the logit model on the basis of instantaneous conditions, and the choices actually being made by travellers at the same point of simulated time. These proportions need not be the same, but we assume that as people learn about conditions and adapt their behaviour, the actual choices will tend towards the ideal. This is a common device in system dynamics models: the logit models deliver the target shares, and the actual shares are gradually adjusted towards them.

Figure 8 illustrates how this is handled in the context of route choice for car trips. The current actual route choice proportions are held as stocks. From these it is possible to calculate the current number of people travelling on each route, and hence on each link in the network. These numbers are then converted into flows (‘trips to work per hour’) simply by dividing by a time period: for travel to work this would typically be two hours, or so.

Figure 8: A dynamic logit model to model choice of driving route.



Each link has a capacity, an ability to handle traffic flow. If the flow rises above this the effect is congestion, and the traffic speeds drop. By comparing the actual flows with the capacities on each link it is possible to calculate an adjustment to the travel times, and hence a total, current, route drive time.

Other terms can be added in at this point to allow for such things as congestion charges. However eventually we reach a point where the utility of each available route between every pair of zones can be calculated, on the basis of current levels of congestion, and a logit model can then calculate the target choice proportions. The actual route choice proportions are adjusted towards these targets in simulated time as people learn and adapt. As they do so of course, they may change the congestion conditions on the road, or crowding on public transport, and hence the target choice proportions.

It is worth noting here that congested route assignment, which this is, is notoriously difficult to handle in traditional equilibrium models because of the feedback loops. Within a system dynamics framework it is relatively simple to formulate and implement.

Implementation

This section outlines some applications of the model that have been made over recent years. The model is, of course, still evolving, and is being adapted as circumstances allow.

An early version of the model, built in Ithink, was used in a study of Hastings, on the south coast of England. Following that the model was completely rebuilt in Vensim, and has now been used in a study of Merseyside (Liverpool) in North West England, and, currently in a major study of the north east of England. A test application based around a hypothetical town has also been constructed.

Hastings

Hastings is a town on the south coast of England. Although the South East is the most prosperous part of the country, Hastings has long been in decline. Its older industries of tourism and fishing have largely gone, and it is too far from London for any significant commuting. It suffers however from severe through traffic passing along the coast, and a by-pass had been proposed for many years, along with improved links to the north, towards London. The objective of this study was to assess the by-pass (plus other transport schemes); a significant sub-objective was to assess the extent to which new road links would help regenerate the town's economy, a claim frequently made by advocates of the schemes. The by-pass would have provided access to a pocket of land on the edge of the town, largely inaccessible now, that could be released for development, and the argument was that this would provide the impetus for the regeneration needed.

Conventional transport models were built, and programs of surveys and interviews with businesses in the town carried out. The first DUM was built to explore the dynamics of what might happen if the roads were built. In summary, the findings were:

- That the roads would increase the range of commuting, both into and out of the town, at least for people with access to cars. This happened because of the reduced congestion delays and the increased speeds on the by-pass, giving car-owning residents of Hastings access to employment opportunities elsewhere. However it also provided new competition for what jobs there were in Hastings, and this put people in the town without a car at increased disadvantage. The total volume of car travel increased, mainly due to increased driving distances, but there was little impact on the levels of economic activity in the town;
- With the land released for development, the model did predict that it would eventually be developed, with houses and business units. New jobs would be created, and commuters would be drawn in to the site from a fairly wide area – not just Hastings. However there was little improvement in economic activity in

the town itself, and in some scenarios jobs there were actually lost as activity transferred to the new site.

It might be added that this scheme was hugely controversial, for it passed through an area of outstanding natural beauty, and the government's position at the time was a presumption against road building in such locations. The 'pro' lobby was very strong and active, and it was certainly the case that the problems due to through-traffic in the town were bad. In the event however the government decided not to build the by-pass. Part of the argument was that the model (and other work) had shown there would be increased business activity on the town edge and increased migration to take up the new housing and jobs on offer, but because of this employment benefits would go to primarily to new migrants to the town, not to existing residents. The objective of raising their welfare was not shown to have been achieved.

There is an interesting footnote to the story. A few months after the decision, it was announced that £125m was to be made available to support a regeneration programme for the town. This was the money that would have been spent on the by-pass: now it will be spent on refurbishment, internet technology and a new University in the town centre.

Hypothetical town

A hypothetical town has been set up largely to test the model as it is developed, but which is capable of demonstrating interesting dynamics. This town consists of seven zones, with simple road and public transport networks linking them. There is no transport link with the outer world.

Figure 9 illustrates the town. It lies on the coast (the design was inspired by Hastings, but it is not that town). The red lines represent the roads linking the zones, and the green lines a small public transport network. The links are directional: the model recognises the distinction between link 1a, for example, which carries traffic in the direction of the arrow shown, and its reverse, which carries traffic in the opposite direction.

Figure 9: Hypothetical town



The model is initialised so that it is in reasonably stable equilibrium, so without some interventions not much changes over ten simulated years. Various tests can then be implemented.

As with many countries, road pricing is being considered in the UK, usually to reduce congestion, and in fact has recently been introduced in London. Increased public transport capacity is usually proposed to accompany the charges (as is the case in London). The Dynamic Urban Model is well suited to test the likely impacts of road and congestion pricing, so the first test described here is to impose a fixed cost of £5.00 on every car with a destination in zone 6, the town's central area.

Figure 10 shows the effect on the total number of jobs in the town compared to the base run: there is a moderate decline, in response to what is effectively an increase in the cost of doing business. Figure 11 shows how this looks in Zone 6, for finance and service businesses: a slight decline, compared to the base, followed by recovery. However it can be seen in Figure 12 that the effect on retail is very severe, with a drastic decline in the number of retail businesses.

This can be compared to Figure 13, which shows the number of retail businesses in Zone 1, on the edge of the town. Here there is an increase, albeit not large numerically. Why does this happen for retail, but not other sectors (the finance and services sector changes very little in Zone 1)?

It is due to a combination of two factors. First, the charge deters retail customers from visiting zone 6, who then switch their spending to other locations. Second, it makes recruitment of a staff harder for companies based on zone 6, and as people switch jobs over time they tend to switch to locations with cheaper access. The result is that

conditions improve for businesses outside Zone 6, because now it is easier to recruit, and they pick up more retail business: eventually the result is the shift shown.

The effect is less severe for finance and service businesses because they are less affected by loss of customers than the retail sector, and in the longer term they also gain from the premises vacated by retail businesses that leave, which they can make use of. This is why after a period the number of finance and service businesses in zone 6 recovers.

Figure 10: Impact on total jobs of £5 fee in Zone 6




Figure 11: Finance & service businesses in Zone 6

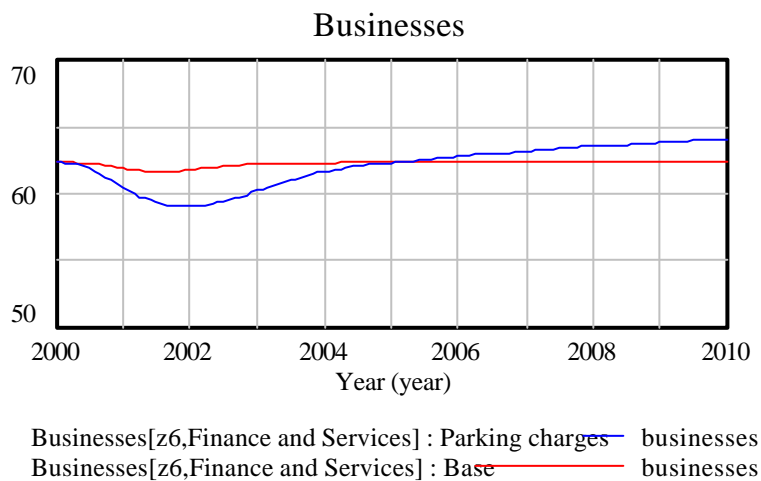


Figure 12: Retail businesses in Zone 6





Figure 13: Retail businesses in Zone 1



The question of mode shift is also of interest, since a shift towards public transport is sometimes an explicit objective of proposed charging schemes. Figure 14 shows the total daily revenue taken by the public transport operator from commuting trips. Compared to the base, the charge generates more revenue for the operator, as might be expected, due to the mode switching of commuters. The revenue rises, but then falls slightly.

Figure 15 plots the change in daily ridership on one of the public transport links into the central area (link 7a). It can be seen that there is an initial transfer to public transport, but this dissipates surprisingly quickly and has virtually gone after four years. The reason is that the people who transfer all have cars and, by and large, prefer to use them. When they next change jobs, the cost of access to jobs outside zone 6 is less, and they tend to switch their place of work, and in so doing move back to car use. In the rather longer term this process is further encouraged by the movement of some businesses out of zone 6 to the rest of the town.

This does not happen on all the links. The zones on the periphery have poorer links with the rest of the town, and so the transfer to other jobs is less easy. The result is that the public transport operator is still left with some revenue gain in the long run.

Figure 14: Total public transport revenue per day


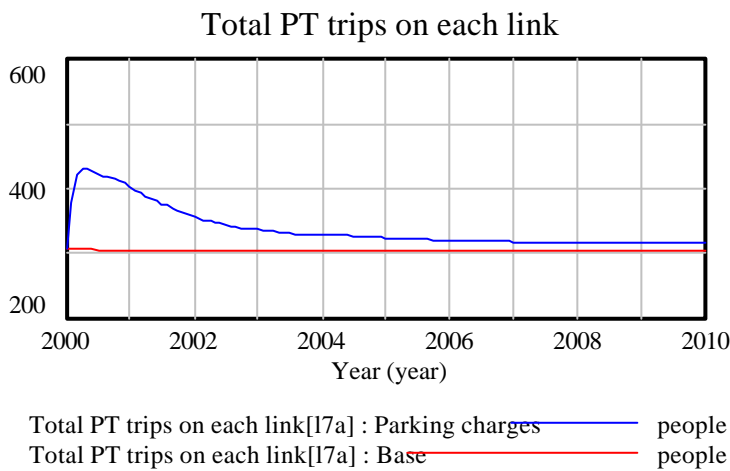


Figure 15: Passengers per day on link 7a



Of course these results are not presented as generalisable to any town. This is after all a rather small test town and £5 is a very high charge given the scale of the place. Moreover the results in any case will depend on the distribution of residences and businesses and the transport networks connecting them. What is claimed however is that that the DUM offers an excellent platform for exploring the likely dynamics of road pricing in an urban area, and for exposing the possibility of some unexpected effects, such as the temporary nature of some of the changes.

Merseyside

This study, still under way at the time of writing, was designed to look at the regeneration impacts of a new light rail scheme in the city. This was a larger application of the model in some ways – it contains 140 zones, for example, and the workforce and jobs were segmented by five types of skill. However the application was also more restricted in that many of the dynamics, including population migration and job creation, were switched off, in order to allow externally imposed numbers of jobs and residents to be used. Zone to zone generalised costs for car, public transport and walk, with and without the proposed light rail scheme, were all provided by external conventional transport models.

The only significant dynamic left operating was the recruitment process, and the model was used essentially to see whether the light rail would provide improved access to employment opportunities for people living in severely deprived quarters of the city. As a demonstration of the Dynamic Urban Model's capabilities it is therefore of rather limited value, except in demonstrating that it is capable of handling many zones. However the following conclusions were drawn:

- For households without a car, located in the regeneration areas served by the light rail, the number of accessible jobs increased by about 25%. The figure for those with a car available was very slightly negative, due to the fact that there was some increase in car travel times due to reduced road capacity for cars, but these people continue to give less weight to public transport in their travel choices;
- This generated a modest shift in the balance of advantage for people with and without a car, causing a modest increase in the number of car-less people finding work;
- The overall increase in employment across the whole city was, however, almost zero. As already noted, the model was not able to generate new jobs as a result of the new transport. However there is always a pool of job vacancies, and it was thought that the light rail might improve recruitment and reduce the size of this pool. This did not occur, primarily because there was an excess of unemployed people over available jobs, so that access to staff was not really a constraint. In other words, such vacancies as there were were due to internal delays in the recruitment processes of employers, and not numbers of accessible job seekers. Improving the number of accessible job seekers did not, therefore, reduce the vacancies, although it did shift the geographic distribution and social mix of those who were in work.

The North East of England

The most recent application, under way at the time of writing, is in the North East of England. The economy here was traditionally based on coal and heavy industry, but most of this has gone. There are prosperous areas, but also very deprived areas in which employment is high, housing and health is poor and all the problems that brings. We were asked to look at how transport strategy might be designed to help stimulate

economic regeneration in the region, possibly coupled with other initiatives such as training.

The model has been set up with 81 zones covering the region. Each zone is initialised with information about:

- Its area, and the area allocated for housing or business use;
- The housing stock, split by five types of house;
- The number of households, split by five 6 types; and the structure of each type of household;
- The numbers of businesses operating in the zone, in each of five categories;
- The stock of business premises, in which the businesses are located.

The area covered by the model has a workforce of roughly 1.1m and, currently, rather fewer than 1m jobs. The model has been set up with a link-based representation of the road network that includes the major roads, but not the smaller ones. This amounts to rather more than 1,000 links. Public transport networks were not built, but zone-to-zone generalised costs were inferred on the basis of travel to work patterns and mode choices as reported in the 1991 census.

The model operates in time periods of quarter years. It was run for 30 simulated years to allow transient instabilities to settle out, and then subsequent test runs started from that end point. The full network version of the model takes 1h 15m to simulate ten years, using a 2 Ghz PC, so a second version of the model was also set up to provide a more rapid tool, using fixed zone to zone generalised costs. This takes about ten minutes to simulate ten years, and can be used to provide quick checks of broad based policies, but is not capable of testing specific transport schemes.

This work is still under way at the time of writing, but the following broad conclusions are emerging:

- While there are some deficiencies in the public transport networks and some known congestion spots on the roads, the transport network in the region is quite good. Consequently improvements to the network, while relieving some of the congestion, does not lead to any very great increase in jobs¹⁰;
- On the other hand the balance of advantage between residents can be shifted when transport costs are reduced. Some areas of high unemployment and low car ownership could be shown to suffer some increase in unemployment if improvements to road transport led to increased external competition for those jobs they did have.

Scenarios were tested in which economic growth was imposed externally, such as might happen if the national economy grew. This was modelled by increasing the attractiveness of the region for new businesses, thereby increasing the start-up rates.

¹⁰ This finding is at least consistent with empirical work carried out elsewhere, which has suggested that in economies with well developed transport networks new schemes will not usually lead to increases in jobs unless transport is clearly acting as a constraint

Growth targets were set of 2% of the existing stock per year, in addition to the start-ups already generated by the model. This suggested the following:

- Congestion on the network rose, leading to congestion and some constraint on job growth. However the growth in businesses also increased the attractiveness of each zone for business start-ups, and this tended to counter the effect of the congestion. In other words businesses were tolerant of the congestion because they still had access to growing numbers of businesses. It could be argued that something of this can be seen in successful cities everywhere;
- In many zones the target 2%pa growth could not be achieved because of other constraints, primarily the availability of suitable premises to house the businesses, which in turn was usually due to land shortage. To a lesser degree the availability of a suitable workforce was also a constraint in some cases, but rather as with Merseyside, the high unemployment levels meant that recruitment was not, on the whole, a constraint. In these cases, relieving the transport system did not help to generate new jobs.

Summary

The relationship between transport and the economy is of great interest and importance, but is still little understood, and there is much disagreement about it. Part of the reason for this is the sheer difficulty of obtaining good empirical data, following an intervention to a transport system, and then of demonstrating convincingly that changes to the economy can indeed be attributed to those interventions.

The topic is important because it is frequently claimed that new transport investment is essential to help regenerate towns and cities, even though the mechanisms by which this regeneration will arise are not always explained.

This paper has provided a brief overview of the problem, and of some of the methods that have been used to model these interactions, frequently equilibrium models. It has been argued that equilibrium models are not well suited to the problem, and that system dynamics provides a more appropriate methodology. A model has been described that builds on ideas borrowed from Forrester's original Urban Dynamics work, but extends it to a fully zonal model with road and public transport networks. Four applications of the model, three 'real' and one hypothetical, have been described.

The model is still 'work in progress' and has more shortcomings than the author would like. However it is argued that it represents a significant step forward in the field, and is genuinely practical tool that can help policy makers understand the dynamics of their town or city.

References

Banister D, Berechman J. 2000. *Transport Investment and Economic Development*. UCL Press, London

Ben-Akiva M, Lerman SR. 1985. *Discrete Choice Analysis*. MIT Press, Cambridge, MA

Forrester J, 1969. *Urban Dynamics*. MIT Press, Cambridge, MA

Forrester N, Private correspondence

Standing Advisory Committee on Trunk Road Assessment (SACTRA). 1999. *Transport and the Economy*. UK Department of Environment, Transport and the Regions, London.