The impact of land use effects in infrastructure CBA – a simulation study

Jonas Eliasson, Stockholm City Transportation Administration

Christian Savemark, KTH Royal Institute of Technology

Joel Franklin, KTH Royal Institute of Technology

Introduction

Infrastructure investments are often evaluated ex ante using transport models. The transport model is used to construct a CBA for the investment by calculating demand and generalized travel costs with and without the investment, and based on that social benefits (consumer surplus, producer surplus etc) are calculated and compared with investment and maintenance costs.

In practice, transport models almost always treat residential location as fixed. (Job location, on the other hand, is often allowed to implicitly vary by using a destination choice model for work trips.) It has been argued that this may constitute a source of serious error or bias, since it ignores the fact that the investment will induce changes in residential location.

In this paper, we explore this issue through a simulation approach. We use a stylized city simulation model to compare CBA results for a large number of infrastructure investments in a simulated city with and without taking changes in residential location into account. In particular, we are interested in whether the optimal selection of investments under a budget constraint changes depending on whether residential location effects are taken into account.

The city model includes endogenous land prices and demand for residential land, heterogeneous preferences and wage offers across residents, and spillover mechanisms which affect wage rates in zones. The model is calibrated to generate realistic travel patterns and demand elasticities. The model simulates the very long run, in which residential location and densities are allowed to vary freely, i.e. there is no dependence on historic land use patterns. This is useful for establishing a limit on how large the errors of ignoring land use changes might be.

To check the robustness of our conclusions, we analyze four different model setups with different networks, location patterns and distributions of the random taste and wage parameters.

Method

The simulation model

The simulated city is divided into zones containing both residents and workplaces. The zones are connected through a transport network. In the simulations, the geography of the city (zones and network) can be generated randomly, which allows us to check the robustness of our results with respect to different geographic configurations. The geography and the network generates travel times t_{ij} and travel costs c_{ij} between each pair of zones (i, j). Travel speeds are constant and equal everywhere, so travel times and travel costs are proportional to link lengths. The population is fixed, and consists of n classes where we treat each class as a continuum. Each class chooses residential zone i, workplace zone j, how many hours to work W and lot size L by solving the following two-stage optimization problem. First, conditional on a residence-workplace zone pair (i, j), the optimal number of working hours W_{ij}^n and the optimal lot size L_{ij}^n are solved for, taking the wage rate offers $\{w_j^n\}$ and the land prices $\{p_i\}$ as given (both the wage rate offers and the land prices are in fact endogenous). Second, the utility of choosing the residence-workplace zone pair (i, j) is assumed to be $\tilde{v}_{ij}^n = v_{ij}^n + \xi_{ij}^n$ where ξ_{ij}^n is a random term capturing idiosyncratic taste preferences, different for each individual in the class and assumed to be redrawn when the scenario changes (e.g. following a transport investment). Individuals in the class choose the residence-workplace pair which maximizes $\tilde{v}_{ij}^n$. This gives aggregate demand for land in each zone. Land prices are then calculated to make demand equal the fixed supply of land in each zone. To emulate external agglomeration effects in the model, the wage rate offer in zone *j*, w_j^n, is assumed to increase with the number of workers in that zone. This emulates spillover effects such as sharing and learning between workers and firms. This obviously also affects the equilibrium distribution of workers and jobs.

The parameters of the model are calibrated to make key outputs roughly correspond to actual values from Stockholm (e.g. wage distribution, average travel time, residential density distribution, elasticities of travel time and output and so on) of the model.

In order to check the robustness of our conclusions, we study four different networks (shown below) with roughly the same number of links, zones and total area: a grid, a star-shape and two randomly generated networks. The overall model parameters are the same in each setup, but the vectors of taste parameters and wage offers are redrawn (from the same distribution)

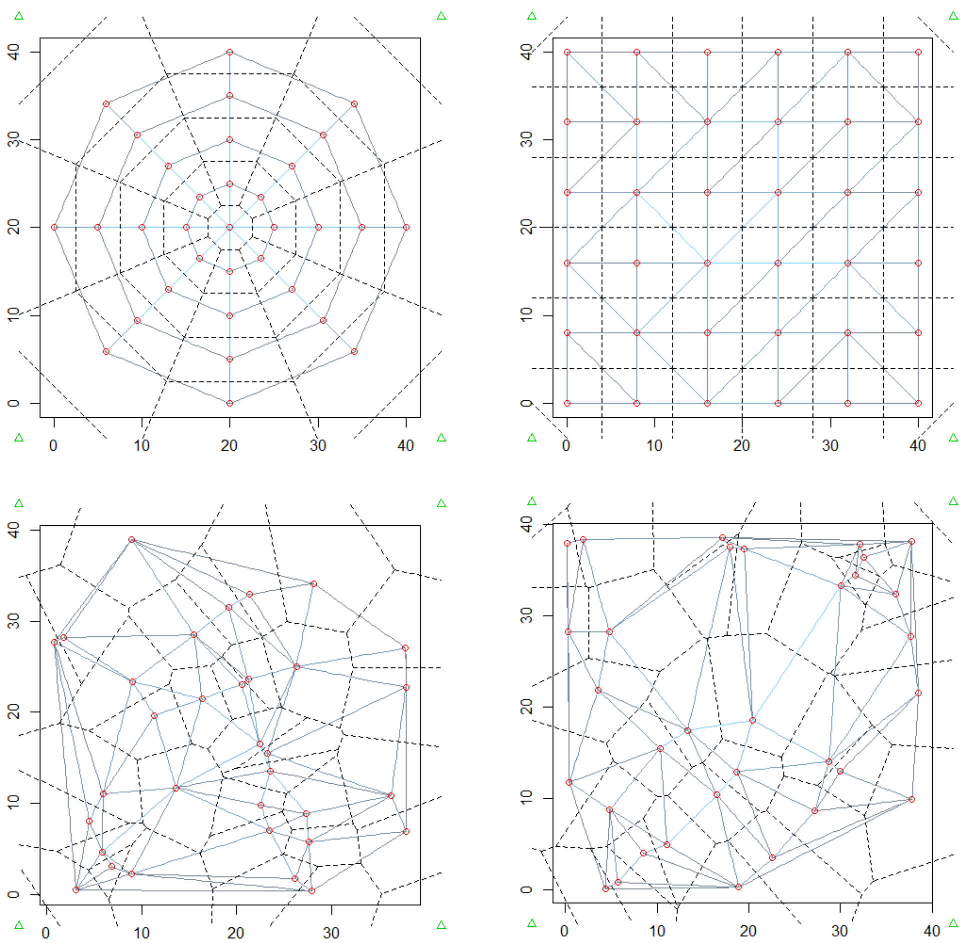

Benefits of a transport improvement

Consider a change in the travel time between an arbitrary pair of zones – say (1,1). In the theoretical part of the paper, we show that the social benefit dB of a travel time change dt_{11} is

$$dB = \sum_{n} P_{11}^{n} \theta_{11}^{n} dt_{11} + \sum_{nij} (1-\tau) P_{ij}^{n} W_{ij}^{n} dw_{j}^{n} - \sum_{nij} P_{ij}^{n} L_{ij}^{n} dp_{i} + \sum_{nij} P_{ij}^{n} L_{ij}^{n} dp_{i} + \tau dZ$$

Here, P_{ij}^{n} is the number of commuters of class *n* from zone *i* to zone *j*, τ is the tax rate, dZ the change in total output, and $\theta_{ij}^{n} = w_{ij}^{n} + \frac{\partial u_{ij}^{n}/\partial t}{\lambda_{ij}^{n}}$ the monetary value of a travel time saving.

The benefits consist of five terms. The first three accrue to workers: travel time savings, i.e. the conventional consumer surplus (the first term), changes in wage rates (the second term) and changes in land prices (the third term). The fourth term is the change in land owner revenues, and the fifth term is the change in tax revenues. Note that the third and fourth terms cancel out: the increase in land owner revenues is simply a transfer from workers to land owners. Hence, if one is interested only in aggregate benefits (and not how they are divided between workers and land owners), one can ignore changes in land prices.

This result is analogous to the result in Eliasson and Fosgerau (2018), who derive a similar formula. The result here is an extension in the sense that the model includes endogenous working hours, worker location, demand for land and land prices.

Results

The central question of the paper is how large error is induced by the conventional modeling practice to assume that residential location is fixed. To explore this, we compare model outcomes where residential location is allowed to vary – which we consider as the "true" outcome – and model outcomes where residential location is fixed – which we consider as the "approximate" outcome corresponding to conventional modeling practice. For each network, we simulate the effects of a 20% travel time reduction on each link in the networks, and calculate the benefits of each such improvement. Each such travel time reduction represents an "infrastructure investment" in that particular link.

The most important comparison is how the prediction error resulting from keeping location fixed affects decisions about what investments to make, i.e. achieved total benefits when selecting the "best" investments under a budget constraint. In addition to total achieved benefits, we also explore how project selection differs, and how large the errors in benefits are.

In this extended abstract, we only present some the results for the grid network (the rest of the results are in the full paper):

Budget constraint	Actual maximal attainable benefits	Benefits when selecting projects using the fixed-location model outcomes	Loss of benefits due to fixed-location assumption in modeling
100	27677.25	26174.19	-5.4%
200	45061.34	42769.03	-5.1%
300	58961.26	56257.26	-4.6%
400	70213.33	65801.93	-6.3%

	Budget 100	*Budget 200*	*Budget 300*	*Budget 500*
Number of projects in "true" optimal selection	12	24	35	57
Number of those projects included in "approximation"-based selection of projects	9	19	29	50
Share of optimal project selection included in "approximation"-based selection	75%	79%	83%	88%

Conclusions

Infrastructure plans are usually constructed, at least in principle, by choosing the best projects from a large pool of candidates, subject to an overall budget constraint. The objective – again, at least in principle – is to attain maximal total benefits given this budget constraint. To do this, policy makers use the output of transport models to assess the effects and benefits of the competing investments. However, transport models used in practice (and hence the CBA:s based on the outputs) almost always neglect changes in residential location. The central question of this paper is how large the error caused by this approximation is.

Broadly speaking, our results indicate that neglecting to account for changes in residential location (and hence changes in demand for land) has only a marginal impact on the attained total benefits of an infrastructure plan – at worst a loss of around 5% of the total attained benefits. However, the difference in project selection is larger: around 20% of the projects that *should* have been included in the project selection are left outside. Still, considering other sources of error in the modeling and decision-making process, this is perhaps not a major concern.