Welfare Analysis in Spatial Equilibrium

Daniel Hörcher[1,2,3]

[1]Senior Research Associate, Corvinus University of Budapest, Hungary
[2]Research Associate, Budapest University of Technology and Economics, Hungary
[3]Honorary Research Fellow, Imperial College London, U.K.

Short summary

Transport cost-benefit analysis (CBA) is a key tool in policy-making, yet its methodology has remained largely unchanged for decades. This paper integrates transport appraisal into a quantitative spatial model (QSM) with endogenous time valuation, enabling welfare analysis in both partial equilibrium (PE) and spatial general equilibrium (SGE) settings. We quantify welfare using two competing approaches: a conventional PE method based on direct user benefits and approximated agglomeration effects, and a structurally consistent SGE framework that captures the full reorganisation of economics activity in space. The latter is implemented via a welfare decomposition that aligns closely with standard CBA components. Empirically, the resulting welfare estimates are of similar magnitude across approaches, with SGE estimates falling within the range of PE outcomes. This suggests that adopting a general equilibrium framework is unlikely to fundamentally alter project rankings. Monte Carlo simulations confirm a strong correlation between methods, with differences emerging primarily for smaller interventions.

Keywords: transport CBA; spatial equilibrium; wider economic impacts

1 Introduction

Transport cost-benefit analysis (CBA), often referred to as benefit-cost analysis (BCA) in the US and transport appraisal in Europe, is one of the most impactful areas of application of microeconomic theory. In the transport sector, welfare economic models are frequently applied in policy-making to prioritise among a pool of proposed infrastructure investments and justify the societal value of allocating public funding to the construction of transformative megaprojects.

Despite its undeniable impact on policy-making, the mainstream methodology of transport appraisal has remained nearly unchanged over the past two decades. This paper contributes by enriching a quantitative spatial model (QSM) with the essential ingredients of transport policy analysis, including endogenous time valuation. It delivers a spatial general equilibrium model that can be quantified using the advanced empirical features of a QSM. Our QSM specification includes the leisure–labour trade-off, which provides microfoundation and a simple but spatially differentiated analytical formula for the money value of time, in line with the canonical literature of transport economics (Small, 2012). As the model reproduces the spatial distribution of travel time valuations, the general equilibrium welfare results can be directly benchmarked against their partial equilibrium counterparts. Using recent theoretical results from Donald et al. (2025), we decompose the aggregate welfare change in spatial equilibrium into additive components that are also directly comparable with the standard elements of a mainstream CBA, including direct user benefits, the agglomeration externality, and financial implications. Furthermore, because model estimation and policy simulation are performed within a single integrated framework, the paper advances the quantification of uncertainty around the core CBA outputs: we compute the standard errors for each component of the aggregate welfare change, as well as for several spatially distributed estimation outcomes.

The quest to produce spatially explicit numerical models that capture the impact of transport policies on location choice, land-use and other socio-economic outcomes has a long history in transport research. Hörcher et al. (2026) review previous research on land-use transport interaction

(LUTI) and spatial computable general equilibrium models (SCGE) from the past half a century. These approaches showcase remarkable progress in simulating complex spatial phenomena, but they have not become part of the consensual CBA toolbox. The absence of a causal and theoretically coherent model estimation method is a frequently cited shortcoming of these early spatial models, which often rely on hundreds of arbitrarily selected exogenous parameters.

Since the mid-2010s, a new stream of research often referred to as *quantitative spatial economics* (Redding & Rossi-Hansberg, 2017) has made major advances in developing spatial equilibrium frameworks that can be used simultaneously for causal parameter estimation and policy evaluation. Quantitative spatial models (QSMs) are SGE models designed specifically to facilitate parameter inversion and the use of causal techniques in model estimation. Despite its popularity in urban economics, the QSM technique is nearly unknown in the international community of transport economics.

This presentation focuses on one specific aspect of the long version of the paper, Hörcher & Graham (2025): how economic welfare measures can be derived from QSMs. We exploit the fact that our model specification yields an endogenous measure of the value of time, which can be derived analytically and quantified using standard spatial data. Beyond direct user benefits, we quantify external firm productivity gains in two ways. First, we approximate the change in economic density following the transport policy and the resulting increase in firm output. Second, we adapt the welfare decomposition technique of Donald et al. (2025) to our setting, thereby capturing the full general equilibrium structure of the model. This benchmarking exercise is an essential step towards making the QSM methodology fully compatible with mainstream transport cost-benefit analysis.

2 Methodology

Model specification

Let us define the utility of a representative worker who resides in location i and commutes to location j as

$$U_{ij} = \left(\frac{L_{ij}}{1-\gamma}\right)^{1-\gamma} \left(\frac{K_{ij}}{\gamma}\right)^{\gamma} z_{ij}; \quad K_{ij} = \left(\frac{C_{ij}}{\beta}\right)^{\beta} \left(\frac{H_{ij}^R}{1-\beta}\right)^{1-\beta}. \tag{1}$$

In this specification L_{ij} is a measure of leisure time, K_{ij} is the composite subutility derived from consumption C_{ij} and residential floorspace use H_{ij}^R, γ and β are structural parameters, and z_{ij} is an idiosyncratic taste shock associated with the combination of locations i and j. To keep our notation simple, we suppress the unique identifier of households. Note, however, that z_{ij} takes a different value for each household.

Commuters are confronted by two constraints. First, wage w_j at workplace j and the monetary price of commuting τ_{ij} times individual labour supply x_{ij} determine the budget available for consumption, given the price of the consumption good, p_i, and the price of residential floorspace, q_i.

$$x_{ij}\left(w_j - \tau_{ij}\right) = p_i C_{ij} + q_i H_{ij}^R \quad [\kappa_{ij}] \tag{2}$$

Second, the sum of leisure time L_{ij} and the total time spent at work (T) and in commute (t_{ij}) cannot exceed $\bar{L}$, the daily time endowment of households.

$$\bar{L} = L_{ij} + x_{ij}\left(T + t_{ij}\right) \quad [\mu_{ij}] \tag{3}$$

Again, individual labour supply $x_{ij} \geq 0$ determines how many workdays an individual is willing to provide.

The idiosyncratic utility shock is specified as an i.i.d. random draw from a Fréchet distribution:

$$F(z_{ij}) = \exp(-X_i E_j z_{ij}^{-\epsilon}), \tag{4}$$

where the average amenity (i.e. the scale parameter) is defined as the product of residence and workplace dependent local fundamentals X_i and E_j, and ϵ governs the spread of individual preferences. These assumptions lead to location choice probabilities that take the form of a commuting

gravity equation (Ahlfeldt et al., 2015, Appendix Section S.2.3),

$$\lambda_{ij} = \frac{X_i E_j \left[\frac{\upsilon_{ij}}{p_i^{\beta} q_i^{1-\beta}}\right]^{\gamma\epsilon}}{\sum_r \sum_s X_r E_s \left[\frac{\upsilon_{rs}}{p_r^{\beta} q_r^{1-\beta}}\right]^{\gamma\epsilon}}, \tag{5}$$

where $\upsilon_{ij} = \mu_{ij}/\kappa_{ij} = (w_j - \tau_{ij})/(T + t_{ij})$ can be interpreted as an effective wage or the marginal value of time defined as the ratio of the marginal utilities of time and money. Later we will use this endogenous measure of time valuation to quantify the benefit of time savings in partial equilibrium welfare analysis.

To model the production side of the economy we follow the conventional approach of quantitative urban modelling more closely. Production in location j is governed by a Cobb-Douglas function of total labour input M_j^W and commercial floorspace H_j^W, with expenditure shares α and $1 - \alpha$.

$$Y_j = A_j (M_j^W)^{\alpha} (H_j^W)^{1-\alpha} \tag{6}$$

We assume perfect competition in the goods market combined with zero trade cost within the urban area. Following the empirical literature of agglomeration (Combes & Gobillon, 2015), we decompose local firm productivity into an exogenous component and a multiplier dependent on a measure of economic density, i.e. access to economic mass (ATEM), denoted by ρ_j.

$$A_j = a_j \, \rho_j^{\eta}; \quad \rho_j = \sum_s \exp(\delta t_{sj}) \, M_s, \tag{7}$$

Parameter η is the agglomeration elasticity of firm productivity, and ρ_j is a measure of access to economic mass (ATEM). M_s is a measure of economic activity (employment), and $\delta < 0$ is the rate by which agglomeration economies decay over travel time t_{ij}. In a separate section of the long version of this paper (Hörcher & Graham, 2025), we estimate $\hat{\eta} = 0.142\,(0.024)$ and $\hat{\delta} = -0.071\,(0.012)$.

Finally, our setup considers a third group of agents too: the construction sector, which turns land and capital into residential and commercial floorspace. The construction sector endogenises land-use and the spatial distribution of housing prices. As this sector does not play a role in the rest of this short paper, we refer the reader to Hörcher & Graham (2025).

Appraisal in partial equilibrium

Using the paper's previous notation, the measure that the literature and practitioners often call *direct user benefit* is an application of the "rule of one-half" in microeconomics:

$$B = N \sum_{ij} \frac{1}{2} \left(\lambda_{ij}^0 x_{ij}^0 + \lambda_{ij}^1 x_{ij}^1\right) \cdot \upsilon_{(ij)} (t_{ij}^1 - t_{ij}^0). \tag{8}$$

This is an approximation of consumer surplus in partial equilibrium, assuming a linear inverse demand curve for travel demand that expresses the generalised travel price in function of travel demand (Glaister, 1974). In equation (8), $0.5\,N(\lambda_{ij}^0 x_{ij}^0 + \lambda_{ij}^1 x_{ij}^1)$ is the mean of the commuting volumes in periods 0 (before the intervention) and 1 (after the intervention), and $\upsilon_{(ij)}(t_{ij}^1 - t_{ij}^0)$ is the monetary value of the travel time gain/loss between the two periods. That is, this measure includes both the time savings of those who travelled before the policy intervention as well as the surplus associated with newly induced demand. Note that our model enables us to differentiate the value of time, $\upsilon_{(ij)}$, between the commuting origin–destination pairs. However, as such differentiation is rarely performed in the state-of-the-practice,[1] we will compute PE results with both homogeneous and heterogeneous time valuations.

In the PE appraisal exercise we would like to consider the agglomeration externality stemming from improved access to economic mass and the associated productivity gain, in line with the transport appraisal literature (Graham & Gibbons, 2019). The common methodology in this area assumes, based on the urban economic theory developed by Venables (2007), that the additional

[1] For an overview of the equity and policy concerns underlying the differentiation of the value of travel time savings, see Börjesson & Eliasson (2019).

output that firms realise through improved ATEM, keeping all input factors constant:

$$\Delta Y_j = Y_j(\rho_j^1) - Y_j(\rho_j^0) = \left[\left(\frac{\rho_j^1}{\rho_j^0} \right)^{\eta} - 1 \right] \cdot Y_j(\rho_j^0). \tag{9}$$

where $t \in \{0, 1\}$ indicates the equilibria before and after the intervention. We compute this quantity for each location j based on the changes in ATEM in response to the transport policy. We then sum the resulting ΔY_j values to obtain an aggregate measure of the agglomeration externality.

An important methodological dilemma we need to address at this stage is whether we consider relocation in the calculation of PE welfare results. Relocation affects both direct user benefits through $\lambda_{ij}^1 x_{ij}^1$ in (8) as well as the external productivity gain through ρ_j^1 in (9). As there is no clear theoretical guidance on whether relocation should be considered in PE welfare analysis, we compute both in the following investigation. In the Results section we refer to the PE models with and without relocation as *dynamic* and *static*, respectively.

Welfare analysis in spatial general equilibrium

The literature provides at least two methods to quantify the change in aggregate welfare in a comparative statics exercise of our SGE model. The first one follows a recent contribution by Donald et al. (2025). Their method allows us to derive the aggregate welfare change as the sum of five separately interpretable components: (i) time savings, (ii) rearrangement of consumption and (iii) residential floorspace use to places with high marginal utility, (iv) fare revenues, (v) landowners' revenues, and (vi) external productivity gains. This section reports the core assumptions and final results of the analytical work we performed following Donald et al. (2025), while the Appendix of Hörcher & Graham (2025) contains the detailed derivations.

We define a pseudo-planning problem whose solution is equivalent to utility-maximising location choices and the corresponding general equilibrium allocation in the spatial model. The Langrangian of this pseudo-planning problem is

$$\begin{aligned}
\mathcal{L} =& \mathcal{W} \left(\sum_{ij} N\lambda_{ij}(\Theta) \cdot u_{ij}(L_{ij}, C_{ij}, H_{ij}^R) - \Psi\left(\{\lambda_{ij}(\Theta)\}\right) \right) \\
&+ \sum_{ij} N\lambda_{ij}(\Theta) \cdot x_{ij}\tau_{ij} + \sum_j p_j^{\ell} \cdot \ell_j \\
&+ \tilde{p} \left[\sum_j A_j f_j^y(M_j^W, H_j^W) - \sum_{ij} N\lambda_{ij}(\Theta) \cdot C_{ij} \right] \\
&+ \sum_j \tilde{w}_j \left[\sum_i N\lambda_{ij}(\Theta) \cdot x_{ij} - M_j^W \right] \\
&+ \sum_j \tilde{q}_j \left[f^h(Z_j, \ell_j) - \left(\sum_k N\lambda_{jk}(\Theta) \cdot H_{jk}^R \right) - H_j^W \right] \\
&+ \sum_j \tilde{p}_j^{\ell} \left[\mathbb{L}_j - \ell_j \right].
\end{aligned} \tag{10}$$

In this problem, the objective includes (i) the $\mathcal{W}$ function of household utility across all ij pairs, (ii) fare revenues given the monetary commuting cost vector $\{\tau_{ij}\}$, and (iii) the land revenue collected by absentee landlords. $\Theta = (\{L_{ij}\}, \{C_{ij}\}, \{H_{ij}^R\})$ is a simplified notation for the spatial distribution of leisure time, consumption and residential floorspace use patterns, i.e. the determinants of location choice. The derivative $\Lambda = \frac{d\mathcal{W}(x)}{dx}$ captures the social value of a marginal increase in household utility. Following Donald et al. (2025), we apply a normalisation that implies

$$\Lambda = \mathbb{E}_{ij} \left[p \left(\frac{\partial u_{ij}}{\partial C_{ij}} \right)^{-1} \right]. \tag{11}$$

That is, we normalise household utility by the expected marginal utility of consumption across locations to express its contribution to money-metric social welfare. $\mathcal{W}$ depends on two additive

components. The first one is the total systematic utility that households realise through local leisure time, consumption and housing quantities. The second one, $\Psi(\{\lambda_{ij}(\Theta)\})$, measures the shift in total utility caused by ex-ante idiosyncratic location preferences under a given distribution of location choices captured by choice probability vector $\{\lambda_{ij}\}$. The remaining parts of equation (10) are constraints associated with four market clearing conditions in the production, labour, construction and land markets, respectively.

Consider a set of transport policy interventions that induce small changes in the travel time matrix $\{t_{ij}\}$. Based on equation (10) and the envelope theorem, the associated change in aggregate welfare is

$$\begin{aligned} dW = & \underbrace{\sum_{ij} \Lambda \mu_{ij} \cdot N_{ij} x_{ij} \cdot (-dt_{ij})}_{\text{(i) time savings}} \\ & + \underbrace{N \cdot \mathrm{Cov}(\Lambda - \kappa_{ij}^{-1},\, \kappa_{ij} p \cdot dC_{ij})}_{\text{(ii) dispersion of MU of consumption}} + \underbrace{N \cdot \mathrm{Cov}(\Lambda - \kappa_{ij}^{-1},\, \kappa_{ij} q_i \cdot dH_{ij}^R)}_{\text{(iii) dispersion of MU of housing}} \\ & + \underbrace{\sum_{ij} \tau_{ij} \cdot d(N_{ij} x_{ij})}_{\text{(iv) fare revenues}} + \underbrace{\sum_j \ell_j \cdot dp_j^{\ell}}_{\text{(v) land value}} + \underbrace{\sum_j \frac{p \eta Y_j}{\rho_j} d\rho_j}_{\text{(vi) agglom. externality}} . \end{aligned} \tag{12}$$

Item (i) transformes time savings into social welfare, using the origin–destination-specific marginal utility of time (μ_{ij}) and the average inverse marginal utility of money that we defined in equation (11). Items (ii) and (iii) originate from the unequal spatial distribution of the marginal utilities (MU) of consumption and housing. The sign and magintude of these terms depend on the covariance of two terms: an ij-pair's marginal utility of money relative to the city-wide average and its change in consumption and housing demand in response to the policy. For example, if dC_{ij} (or dH_{ij}^R) is positive in those residence-workplace pairs where the marginal utility of money is also higher than average, then items (ii) and (iii) are expected to make a positive contribution to the aggregate welfare change. Component (iv) captures fare revenues; even though this cash flow is a transfer within society, it must be taken into account with a positive sign because commuter utility is expressed net of the fare payment. Item (v) is land value uplift; this effect is part of the welfare change because land revenues are retained by absentee landlords as an isolate part of social surplus. Finally, item (vi) is the externality that firms generate when they become better connected and thus more productive.

3 Results

The infrastructure policy we selected as a case study is the implementation of the Elizabeth Line, a large-scale urban rail investment programme that attracted a lot of attention in the academic literature as well. The counterfactual application of the model assumes that the Elizabeth Line is introduced in the baseline 2011 spatial equilibrium of Greater London; we let the model to converge to a new spatial equilibrium, and assess the economic impact of the new intervention in a comparative statics exercise.

First, the left-hand side of Figure 1 shows the change in ATEM in response to the new travel time matrix t_{sj} only, keeping the distribution of employment N_s^W and individual labour supply $\bar{x}_s$ constant. Let us call this the *static* change in access to economic mass. As the travel time effect of the transport project is strictly negative, the static change in ATEM is always positive. In line with intuition, this effect is the greatest near the new railway stations, especially in the suburban parts of the line which may have been more poorly connected before the intervention.

The right-hand side of Figure 1 plots the *dynamic* change in access to economic mass. In this case we allow all endogenous variables, including prices and quantities in the labour and housing markets, to readjust to the new spatial equilibrium. That is, we allow both the masses and spatial impedances to readjust. The primary insight from the plot is that the dynamic change in ATEM is indeed not positive everywhere. In fact, a significant part of the city, including the suburbs in the North and South as well as the Centre of London experience a reduction in economic density. Note from the legend of this figure though that the positive changes in ATEM are much greater in magnitude than the negative changes. Also, the positive changes are highly skewed due to a few

outliers that experience extremely high gains in ATEM.

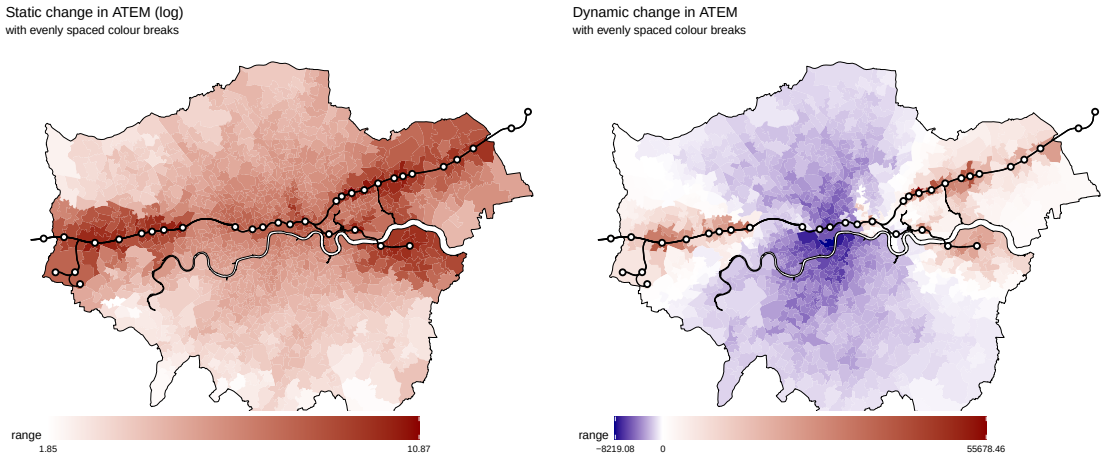


Figure 1: Counterfactual outcomes: static and dynamic changes in economic density (access to economic mass)

Table 1: Benchmarking partial and general equilibrium appraisal methods

Appraisal model	Static partial equilibrium		Dynamic partial equilbrium		SGE
Value of time	homogeneous	heterogeneous	homogeneous	heterogeneous	n.a.
Direct user benefit	830.97***	752.84***	1082.4***	940.81***	810.04***
	(1.66)	(1.4)	(72.18)	(42.89)	(1.55)
Productivity gain	1,766.59***	1,766.59***	837.20	837.20	1,342.07***
	(669.57)	(669.57)	(547.43)	(547.43)	(498.07)
MUoC dispersion					167.90***
					(59.15)
MUoH dispersion					-21.71**
					(10.84)
Land value uplift					34.61
					(53.11)
Fare revenues	1.33	1.33	1.33	1.33	1.33
	(26.27)	(26.27)	(26.27)	(26.27)	(26.27)
Total benefit	2,598.88***	2,520.75***	1,920.93***	1,779.34***	2,334.23***
	(647.04)	(512.06)	(647.06)	(524.68)	(498.73)

Standard errors in parentheses, ***: 99%, **: 95%, *: 90%
Units: Surplus in thousands of GBP per day.

The results of our core numerical experiment are presented in Table 1. The table reports both the aggregate welfare change (final row) and an additive disaggregation of its components. We also make use of the fact that parameter estimation is an integral part of the appraisal model. This allows us to quantify uncertainty at the level of each welfare component by bootstrapping the entire sequence of estimation and simulation stages and deriving standard errors and significance levels numerically. Uncertainty quantification substantially improves our understanding of the confidence we can assign to various sources of economic gains.

The final column of Table 1 presents the welfare measure of the SGE model. In this case, the split between user benefits and productivity gains remains similar in magnitude to the static PE model, although both quantities are somewhat lower. Although the productivity externality is also significant even at the 99% level, statistical uncertainty is considerably lower for the estimated direct user benefit. At the lower end of the 95% confidence band around the productivity gain, the estimate falls to less than half of the direct user benefit.

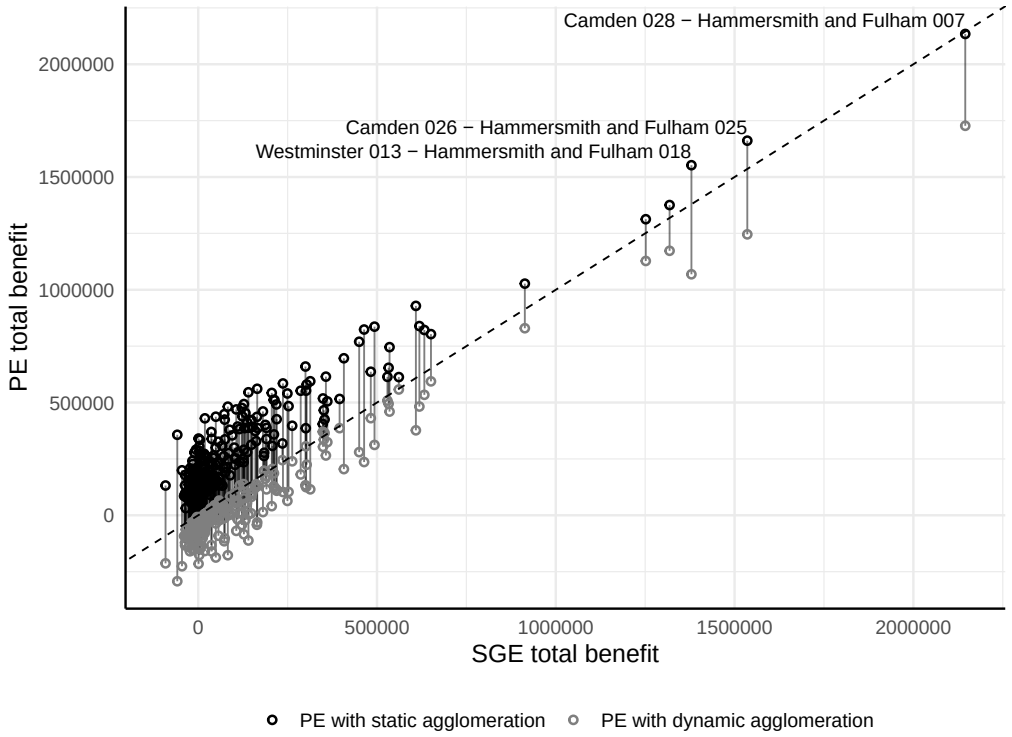


Figure 2: Welfare results for 250 randomly generated transport improvements using the partial equilibrium (PE) and spatial general equilibrium (SGE) methods. The dashed line is the 45° ray along which the PE results coincide with the SGE results.

4 Conclusions

Our core insight is that, although the SGE model may seem methodologically fundamentally different from the more common PE approach, the two sets of aggregate welfare estimates remain of the same order of magnitude. All five estimates of the total benefit are statistically significant, and the SGE model has the second highest t-statistic. This is reassuring from a practical perspective: in general, it is unlikely that the SGE model would overturn many past policy interventions that were justified using the mainstream PE CBA. Under this paper's assumptions on model specification, the SGE welfare estimate is around 10% lower than the highest PE estimate and 31% higher than the lowest. This implies that the SGE model is not only more informative as a disaggregate spatial impact assessment tool, but also a source of evidence suggesting that the magnitude of bias, caused by double counting and other methodological inconsistencies in PE appraisal, may lead to only modest deviations in PE benefit-to-cost ratios.

To make further inferences about the external validity of the paper's core findings, in particular the magnitude of benefits in SGE relative to the static and dynamic approaches of mainstream PE appraisal, we performed a randomised experiment with a series of hypothetical transport improvements. Each Monte Carlo scenario assumes travel time reduction between a randomly selected central MSOA and a peripheral zone. The core results are shown in Figure 2. We observe that across 250 randomly constructed counterfactual episodes, the high correlation between PE and SGE appraisal benefits remains remarkably robust. In relative terms, small improvements are more susceptible to diverging appraisal estimates. For example, when the SGE model predicts zero or even negative welfare effects, the static PE method may still indicate non-negligible gains. The tendency of the PE approach with inelastic location choice to produce higher benefits also emerges as a consistent pattern.

References

Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., & Wolf, N. (2015). The economics of density: Evidence from the Berlin Wall. *Econometrica*, *83*(6), 2127–2189.

Börjesson, M., & Eliasson, J. (2019). Should values of time be differentiated? *Transport Reviews*, *39*(3), 357–375.

Combes, P.-P., & Gobillon, L. (2015). The empirics of agglomeration economies. In *Handbook of regional and urban economics* (Vol. 5, pp. 247–348). Elsevier.

Donald, E., Fukui, M., & Miyauchi, Y. (2025). *Unpacking aggregate welfare in a spatial economy.* (NBER Working Paper No. w34075)

Glaister, S. (1974). Generalised consumer surplus and public transport pricing. *The Economic Journal*, *84*(336), 849–867.

Graham, D. J., & Gibbons, S. (2019). Quantifying Wider Economic Impacts of agglomeration for transport appraisal: existing evidence and future directions. *Economics of Transportation*, *19*, 100121.

Heblich, S., Redding, S. J., & Sturm, D. M. (2020). The making of the modern metropolis: evidence from london. *The Quarterly Journal of Economics*, *135*(4), 2059–2133.

Hörcher, D., Cusson, T., & Graham, D. J. (2026). *Transport in quantitative spatial economics: a review of theory, empirics, and policy simulations.* (Working Paper available at SSRN #6155346)

Hörcher, D., & Graham, D. J. (2025). *A quantitative urban model for transport appraisal.* (Working Paper available at SSRN #5067099)

Redding, S. J., & Rossi-Hansberg, E. (2017). Quantitative spatial economics. *Annual Review of Economics*, *9*, 21–58.

Small, K. A. (2012). Valuation of travel time. *Economics of Transportation*, *1*(1-2), 2–14.

Venables, A. J. (2007). Evaluating urban transport improvements: Cost–benefit analysis in the presence of agglomeration and income taxation. *Journal of Transport Economics and Policy*, *41*(2), 173–188.