Transport appraisal and quantitative spatial economics: A review of theory, empirics, and transport applications

Report prepared for the Department for Transport

Prof Dan Graham and Dr Daniel Hörcher
Imperial College London, London, SW7 2AZ, UK.
Email: d.j.graham@imperial.ac.uk

January 16, 2024

1 Introduction

1.1 Background

The official methodologies of transport appraisal in the UK and other countries worldwide have evolved from a partial equilibrium approach. Such an approach considers the transport market in isolation, ignoring the impact of transport interventions on the spatial distribution of economic activity and market failures in other markets. The limitations of the partial equilibrium approach have been recognised for decades. For example, in a discussion paper of the Department for Transport, the authors state the following:

> *"[The limitations of static models] suggest that a more rigorous approach, such as general equilibrium (GE) modelling, could deliver more robust estimates. [...] A GE approach requires a much more complex model and significantly more data. So far there are few GE models that can credibly model changes in transport costs, let alone provide estimates for particular schemes. Lack of data crucial in GE modelling provides a further barrier to applying such a framework in the UK. Until there are further developments in GE modelling for the UK, we advise that estimates of this wider benefit are based on the partial analysis."*
> (Department for Transport 2005)

Recognising the importance of this knowledge gap, there have been many attempts to produce spatial models suitable to predict the impact of transport policies on urban form and robustly quantify welfare effects in general equilibrium. Some of these attempts, including land-use/transport interaction (LUTI) models (Hunt and Simmonds 1993, Acheampong and Silva 2015) and spatial general equilibrium models (SCGE, see a review in Robson et al. 2018) have been successfully applied to measure "dynamic" effects in appraisal, acknowledging that land use changes as a result of the intervention (and other forecast trends). However, none of these modelling approaches has reached the necessary level of robustness and academic consensus to replace partial equilibrium cost-benefit analysis (CBA) in the official economic evaluation of transport projects. Specifically, sensitivity with respect to dubiously selected parameter values is often mentioned as a weak point of these models (Tavasszy et al. 2011), while the lack of validation of predicted outcomes is also an often cited barrier.

This report reviews recent developments in the spatial economics literature that could be useful in developing a robust and rigorous general equilibrium approach for transport appraisal. Since the mid-2010s, a new type of model has been proposed in the leading general-purpose economic journals, generating considerable advancements in the field of spatial economics. Redding and Rossi-Hansberg (2017) label the new literature *quantitative spatial economics* (QSE). QSE integrates two traditions – urban economics and economic geography – in one framework. By realising a fully spatially explicit economic model that can be estimated causally, QSE delivers the "technical trick" that Fujita et al. (1999) envisioned in their influential book:

"We would clearly like to be able to carry out such exercises for economic geography—to develop, if you like, 'computable geographical equilibrium' models. [...] Such modeling is not easy. Probably it will be necessary to introduce some new technical tricks to make the models consistent with the data."

1.2 Structure and features of a QSM

The report describes the broad structure and distinctive features of QSMs. We find that the following properties of QSMs are especially important to understand the power of this modelling technique.

- QSMs are a subset of SCGE models. They are micro-founded, that is, their modelling assumptions are fully consistent with microeconomic theory: households are utility-maximisers, firms maximise profits, the actual profits depend on market structure (competition), and the prices and quantities of goods and production factors are determined via market clearing equilibrium conditions. LUTI models do not feature all these properties.
- Just like in other SCGE models, transport plays a key role as a determinant of spatial outcomes: households' residential and workplace choices depend on commuting costs while firms' location choices depend on access to workers, other production factors, and consumers.
- QSMs have two types of parameters.
  - *Local fundamentals* are location-specific measures of geographical characteristics. For example, the inherent attractiveness of the locations for residents and workers, and the magnitude of firm productivity are common local fundamentals in all QSMs. Further local fundamentals are often introduced depending on specific study objectives. Each local fundamental is a vector of values associated with the discrete locations in the spatial model.
  - Spatially undifferentiated *generic parameters* (or *structural parameters*) take the same value across all locations. Examples include the expenditure share parameters in the Cobb-Douglas household utility and production functions, and the agglomeration elasticities and decay parameters in local productivity and amenity levels.
- Both groups of parameters of QSMs can be quantified significantly more robustly than other SCGE models that rely on ad-hoc parameter selection and/or numerical search based on ad-hoc calibration rules.

- The structure of QSMs enables the analyst to find the value of the local fundamentals associated with each location in the model via *model inversion*. What is model inversion? In general equilibrium models economic outcomes such as wages, housing prices and transport costs are derived from known parameters. By model inversion we mean that we recover some of the parameters by assuming that the observed economic outcomes are the solutions of the general equilibrium model.
- Most of the generic parameters of these models can be "causally" identified with econometric tools by directly applying equations in the theoretical model as empirical estimating equations. This enables the analyst to tackle statistical endogeneity concerns via state-of-the-art econometric methods.

The wider class of SCGE models are not suitable for model inversion and causal parameter estimation.

- The production and floorspace sectors of QSMs can be specified either as an urban model with agglomeration economies, or an economic geography model with imperfect competition, or both. It is practical to determine the spatial scale of the model in light of type of the investigated scheme, e.g. intra-urban, regional, or national projects.
- In many QSM studies it is analytically proven that the spatial equilibrium is unique, that is, a policy intervention can lead to only one specific outcome according to the model.
- QSM papers derive a theoretically coherent measure of social welfare through which the economic impact of policy interventions can be assessed. In other words, the QSM method is suitable for economic appraisal.

We conclude that this collection of analytical and quantitative properties are unique to QSMs and thus they materialise a new contribution to the spatial economics literature.

1.3 Main findings

The aim of this report is to review the emerging QSE literature from the viewpoint of transport analysis. We assess whether QSMs can be applied for a transport appraisal methodology. Our core conclusions are the following.

- We identify a trade-off in spatial modelling between the richness of theoretical features capturing the economics underpinning the model, and the robustness of the empirical evidence underpinning parameter calibration/estimation. In this spectrum, QSMs strike a healthy balance between theoretical detail/coherence and reliable parameter estimation. Essentially, they can achieve identification, theoretically and empirically, much more readily than previous LUTI and SCGE approaches.
- Existing QSMs were created by spatial economists for publication in general-purpose economic journals, without any consideration for their application in transport appraisal. Furthermore, the transport community has not adopted these model yet to analyse problems in transport economics. Thus, QSMs are currently somewhat isolated from mainstream transport economics, and never formally applied in the context of transport appraisal.
- Current representations of the transport sectors in published QSMs are overly simplistic. This limitation represents a serious hindrance to their direct adaptation in transport appraisal. In particular, the *iceberg* specification of travel disutility is not consistent with transport economic theory, non-commuting trips are mostly ignored, and route and

mode choice under congestion is inadequately modelled. However, we conclude in the report that all these limitations can likely be addressed via appropriate amendments to the model.

- The QSM framework captures direct user benefits, and the wider economic impacts currently considered in TAG, in an integrated fashion, thus neutralising the double counting concerns that cannot be fully disregarded in an extended partial equilibrium model. It is possible to derive Level 1, 2 and 3 performance metrics from one QSM. However, the general equilibrium framework may not be suitable to partition categories of benefit to the extent we currently do in partial equilibrium CBA. For this reason, the QSM general equilibrium approach would likely complement, rather than replace, the existing partial equilibrium framework.
- The predictive performance of QSMs would have to be rigorously demonstrated before they could be put to use for policy purposes. While these models are theoretically appealing, and encompass a good strategy for empirical implementation, to our knowledge, their predictable power has not been comprehensiveness and rigorously tested. We therefore identify two issues that must be addressed in developing QSMs for appraisal: model validation and uncertainty quantification.
- The report outlines a comprehensive agenda for future research with the ultimate aim of turning QSE into a transport appraisal methodology.

The report is structured as follows. In Section 2 we define what a quantitative spatial model is and outline the distinctive features it has compared to previous spatial-transport models. Section 3 summarises the theoretical underpinnings of this technique. In Section 4 we review the empirical methods behind QSMs and the data requirements. Then we enlist and evaluate the existing transport-related applications in Section 5 before concluding the report in Section 6.

2 What is a Quantitative Spatial Model?

2.1 Terminology

There is an ongoing discussion within the spatial economics community about the definition and use of the terms *quantitative spatial economics* and *quantitative spatial modelling*. This terminology is particularly confusing for a non-specialist audience as the words *quantitative* and *spatial* could refer to a much broader set of models. These adjectives could be used to describe any other models in which spatial characteristics are numerically described. For example, SCGE models also rely on and produce quantitative information in a spatial context, and, therefore, one could argue that these are also quantitative spatial models.

It is important to emphasise that the specialist literature refers to a very specific group of models published since the mid-2010s as *quantitative spatial models* (QSMs). The first contributions such as Ahlfeldt et al. (2015) did not use this expression specifically, but they described their approach, somewhat paradoxically, as a "quantitative theoretical model". Subsequently, Redding and Rossi-Hansberg (2017) titled their early review article of the field "Quantitative Spatial Economics" which is likely the main reason why quantitative spatial modelling and the QSM abbreviation became widely used expressions. This terminology is not exclusive, however. For example, Monte et al. (2018) used "quantitative general equilibrium model" while "quantitative urban model" is also frequently adopted since Heblich et al. (2020). Following these examples, Nagy (2022) names QSMs with monopolistic competition and trade costs a "quantitative economic geography model".

In this report we distinguish QSMs from SCGE models on the basis of a set of distinct properties mostly relating to parameter estimation. Otherwise QSMs are also spatial general equilibrium models that can be used as a simulator of computable counterfactual outcomes. Therefore, we treat QSMs as a subset of SCGE models.

2.2 Distinctive features

Quantitative refers to a specific type of spatial general equilibrium models that can potentially be parameterised or quantified with a high level of precision. QSMs have unique properties that enable model quantification in two ways:

1. A one-to-one mapping exists between observed data (e.g. the commuting matrix, wages, housing prices) and quantitative measures of the fundamental geographic characteristics of each location in the model, including workplace and residential amenities, and firm productivity levels.
2. Most of the remaining generic parameters of the model can be estimated with econometric tools by using equilibrium conditions after algebraic transformations as regression equations or moment conditions.

These properties imply several advantageous features that were not available in previous land-use and SCGE models. First, we are able to measure geographical properties such as the value of "leafy streets and scenic views" (Redding and Rossi-Hansberg 2017) at a high level of spatial granularity. These subjective properties are very difficult to measure with regular empirical methods, e.g. a logit-based residential choice model or a hedonic real estate price model. Geographical fundamentals can thus be quantified practically anywhere at any spatial resolution without looking more closely at the physical determinants of the attractiveness of each location. The ease of building spatially explicit models of regional and urban economies has reduced the attractiveness of stylised-space core-periphery and monocentric city models previously dominating the literature.

Second, the majority of the model parameters are theoretically consistent estimates derived from local data. That is, there is no (or just limited) need to introduce parameters from other empirical contexts. Other general equilibrium approaches are typically based on somewhat different functional forms, with data inputs often collected from other geographical locations and time periods. This problem of inconsistency is significantly reduced in a QSM because the parameters are mostly estimated from data on local economic outcomes which are required for spatial equilibrium modelling anyway, and the empirical functional forms are equivalent to the equilibrium conditions of the simulation model. For example, in the current TAG practice, agglomeration benefits are computed using elasticities estimated from nationwide data in a sophisticated empirical specification. Then the elasticities are applied in a local context and the partial equilibrium CBA assumptions. This inconsistency in geographical scope and functional forms is eliminated in a QSM.

This property is not universal as a small number of generic parameters are still borrowed from other sources. But QSMs operate with very few structural parameters, usually around eight to twelve, and a consistent estimation method is available for the majority of them.[1] One could argue that some of the the structural parameters should differ between locations and the Cobb-Douglas and CES are too restrictive functional forms in the utility and production functions. By accepting the potential cost of compactness, QSMs deliver results based on coherently estimated structural parameters as opposed to past SCGE models which often require hundreds of parameters to be specified, thus giving high rise to a degree of uncertainty and opaqueness.

Third, partly relying on the one-to-one link between geographical fundamentals and the observed data, QSMs have analytically proven equilibrium properties. That is, the existence, stability and even uniqueness of the spatial equilibrium can be proven analytically using the linear algebra approach described in the repeatedly updated but so far unpublished working paper of Allen et al. (2020). This is clearly another distinctive feature as in previous land-use and SCGE models only numerical testing was available to explore equilibrium properties. The analytical proofs ensure more robustness in terms of the counterfactual predictions of QSMs, i.e. that the spatial outcome of a future policy intervention remains unambiguous.

These unprecedented properties come at a price in terms of model flexibility. Certain assumptions and functional forms must remain unchanged to maintain these properties. The way heterogeneity is modelled in QSMs has not changed since the mid-2010s. Quantitative urban models such as Ahlfeldt et al. (2015) assume that households have idiosyncratic preferences associated with each residence–workplace combination in the model, and that these preferences are Fréchet (extreme value type II) distributed.[2] Some of the QSMs with a regional or international scope such as Allen and Arkolakis (2014) assume that firms are heterogeneous in productivity and, once again, their productivity shifters are Fréchet distributed. Due to this commonality QSMs are often called informally “Fréchet models”.

Another common property is that commuting and trade costs are modelled according to the *iceberg* assumption. This approach, borrowed from the international trade and new economic geography literature, assumes that the disutility or monetary cost of transport is a given fraction of the traveller’s baseline utility or the mill price of the traded good. (The *iceberg* terminology comes from the narrative that a part of utility or the traded good melts away along the way from the origin to the destination.) The purpose of this assumption is to keep all utility and production functions multiplicative in the model, thus making functional forms suitable to achieve the advantageous model features above. Arguably, the iceberg assumption violates some of the established practices in transport modelling.[3]

The analytical and quantitative properties described above and the unique model specification allow us to clearly distinguish QSMs from other spatial models. This is also a potential source of criticism, however: QSMs have been criticised recently by Proost and Thisse (2019) for “the

[1] In the simplest case household utility as well as the production and floorspace construction sectors are Cobb-Douglas or CES (with 3 separate elasticities), there is a Fréchet shape parameter and a congestion elasticity (2 more parameters), and agglomeration in firm productivity and local amenities are captured via 2 elasticities and 2 decay parameters. This leaves us with nine exogenous structural parameters.

[2] As Section 3.1 explains in greater depth, QSMs rely extensively on random utility discrete choice theory. In other words, they share common theoretical foundations with empirical travel demand modelling and with mode- and route choice in transport assignment. In general, additive utility functions lead to convenient choice probability expressions with extreme value type I random shocks, while multiplicate ones are similarly easy to handle when combined with extreme value type II; see also (Fosgerau and Bierlaire 2009). As the utility functions in QSMs include a Cobb-Douglas component which is multiplicative, the Fréchet assumption is sensible.

[3] See our discussion on possible ways of overcoming the iceberg assumption in Section 3.4 below.

repeated use of the same functional form for preferences, so that we do not know how robust the predictions found in the literature are," also adding that "at a time of rapid advances in computational power and numerical methods, robustness checks about functional forms are called for."

2.3 Roots in the literature

Quantitative spatial economics and the desire to produce spatially explicit versions of stylised-space economic models have been in the air for decades. We review the roots of this literature in two steps: first, we focus on international trade and then move on to urban economics.

In the context of international trade, QSMs bridge the gap between empirical and theoretical branches of the literature. Empirical models of spatial economic processes relied extensively on various forms of the *gravity equation* since Jan Tinbergen's seminal works in the 1960s to model international and domestic trade flows. Even though gravity equations did achieve considerable success in explanatory power, their early versions lacked proper microeconomic justification, and their functional forms were based on intuition instead of theory. Head and Mayer (2014) provide an extensive overview of the gravity literature. Commuting and trade flows are expressed in a gravity form in QSMs and parameter estimation borrows established techniques from the trade literature.

On the theoretical side, advances in modelling monopolistic competition (Dixit and Stiglitz) and the new economic geography (NEG) approach form the foundations of our understanding of centrifugal and centripetal forces in the spatial economy (Krugman 1991, Fujita and Krugman 2004). There was a genuine need to turn stylised-space NEG and subsequent international trade models into geographically explicit, quantifiable structures. In pursuit of this goal, a number of valuable contributions were made prior to the QSM revolution. Helpman (1998) and Eaton and Kortum (2002) are often cited sources of motivation for QSM authors, in particular in the use of extreme value distributed idiosyncratic firm characteristics. This literature culminated in Allen and Arkolakis (2014), the first quantitative economic geography model.

Concurrently, the urban economics subdiscipline faced a similar challenge: stylised-space monocentric and linear city models produced deep insights into the mechanisms of city formation and some of these insights could be confirmed empirically, but theory could not be turned into a spatially explicit model for a while. Lucas and Rossi-Hansberg (2002), a linear city model, is frequently cited as one of the origins of more recent quantitative urban models although the distinctive features outlined above were not present in the 2002 study. A series of studies by Alex Anas and co-authors deserve special attention in this review. Their urban models show many similarities with today's quantitative urban models in the sense that they were the first to adopt a random utility discrete location choice specification in a general equilibrium model, and their model also includes constants that capture the "inherent attractiveness" of discrete locations.[4] Anas and Kim (1996) and Anas and Xu (1999) are stylised-space urban equilibrium models of a linear polycentric city. Anas and Liu (2007) turned this approach spatially explicit, replicating real geography, and combined the urban equilibrium with stochastic traffic assignment to capture mode and route choice robustly. Later on, Alex Anas and co-authors

[4] Interestingly, Ahlfeldt et al. (2015) may had not been aware of these parallels with the Anas oeuvre as they claim their random utility location choice specification was inspired by the original statistical theory established by Daniel McFadden (1978), and, more specifically, the Eaton and Kortum (2002) approach to modelling heterogeneity in productivity in the trade literature.

developed further implementations of this model for various case study areas including Los Angeles and Paris (Anas 2020, Anas and Chang 2023).

A comprehensive review of spatial economics would describe the literature of land-use and spatial computable general equilibrium models beyond Anas as well. Land use/transport interaction (LUTI) models, including Hunt and Simmonds (1993), make the spatial impact of transport interventions predictable for a selected geography, following the footpaths of Lowry (1964). The SCGE tradition includes a series of attempts to model the spatial distribution of economic activity with thorough microeconomic foundations. Extensive reviews are available in Paulley and Webster (1991), Wegener (2011) and Robson et al. (2018). Here we do not review this literature extensively because, objectively, it is completely ignored in QSM studies, thus implicitly declaring that QSMs have no roots in the LUTI and SCGE traditions. LUTI and SCGE models clearly do not have the distinctive quantitative features enlisted in Section 2.2. Nevertheless, we find it unfortunate that the QSM community does not attach value to LUTI and SCGE models from which many useful ideas could be utilised in their quantitative framework as well.

It is difficult to benchmark LUTI and SCGE models against QSMs objectively. There is an evident trade-off in general equilibrium modelling between the level of detail in the theoretical structure and the robustness of parameter calibration/estimation. QSMs are more parsimonious and thus directly estimatable. SCGE researchers take a different stand in this trade-off prioritising model realism. The cost of their choice is that they need to calibrate hundreds of parameters without consistent sources of empirical evidence. LUTI modellers go even further by allowing certain model components not to have microeconomic foundations at all. While all three strands of the literature achieved considerable success, it seems in recent years QSMs have attracted significantly more attention and appreciation in the global community of spatial economists.[5]

3 Theory and model mechanisms

Quantitative spatial modellers split geographical space into a set of discrete locations (e.g. zones) that represent an area with homogeneous socio-economic characteristics.[6] These nodes of residential and production activities are connected by the transport network represented by further transport nodes and edges between them. Locations and transport links are characterised by a variety of attributes: locations differ in population, economic output, wages, land and real estate prices, while transport links are described by their capacity, vehicle and passenger occupancy, free-flow and congested travel times, etc. Spatial and transport data are transformed into the necessary format and resolution using the rapidly growing collection of geospatial data manipulation algorithms. Most QSM researchers do not apply off-the-shelf commercial software packages. They rely instead on their own code written in one of the open-source computing environments such as Python or R, sometimes making the code publicly available. The purpose of the proposed spatial economic modelling reviewed in this section is to infer unobserved node and link attributes, and compute the value of endogenous attributes

[5] To the best of our knowledge the QSM method has not been in use in local or national policy-making so far but economists at the World Bank and other multinational development banks are using the method increasingly.

[6] Some exceptions, including Allen and Arkolakis (2014), consider continuous geographical space in both economic activities and movement/transport between dimensionless locations.

and infer how they change in response to exogenous socioeconomic shocks or potential policy interventions.

3.1 Random utility location choice

We begin this review with a brief description of household behaviour. Residential and workplace location choice is based on random utility discrete choice modelling (RUM), a methodology well-known in the transport domain. However, the specification of utility to be maximised by the representative household somewhat differs from regular mode and route choice models. The following is a typical specification of *direct utility* in the QSM household problem:

$$u_{ij} = A_i^r A_j^w \cdot f(C_i, H_i) \cdot d_{ij}^{-1} \cdot z_{ij}. \tag{1}$$

Note first that this utility function is multiplicative as opposed to the additive structure of logit-based discrete choice models. Subscripts i and j are residence and workplace indices, and u_{ij} measures the utility associated with a given residence–workplace combination. A_i^r and A_j^w capture the fixed amenity levels offered by i and j. In RUM language, these would be called alternative-specific constants, but this expression is never used in QSM papers. The second key property is that (1) includes an inner utility function $f(C_i, H_i)$ of consumption C_i and floorspace use H_i at residential location i. We will soon return to this part of the household problem. The disutility of commuting between i and j is captured by d_{ij}. This *ad-valorem* or *iceberg* specification implies that commuting disutility is proportional to the remaining determinants of utility; in Section 3.4 we look at this assumption more closely.

Finally, z_{ij} measures the random (idiosyncratic) utility associated with residence–workplace combination ij. The usual assumption is that this unique utility shock is Fréchet distributed with cumulative density function $\phi(z_{ij}) = \exp(-S_i^r S_j^w \cdot z_{ij}^{-\varepsilon})$. S_i^r and S_j^w are *scale parameters* playing a very similar role to A_i^r and A_j^w, capturing location-specific sources of utility. To avoid redundancy, QSMs normalise one of the two pairs of parameters to one. By contrast, *shape parameter* ε has a crucial role in the model as it governs the spread of idiosyncratic preferences for specific locations. The unbiased estimation of these parameters is one of the key objectives of the empirical part of quantitative spatial modelling.[7]

The presence of $f(C_i, H_i)$ implies a discrete-continuous decision process for households. For a given ij location pair, households maximise utility with respect to consumption and floorspace use subject to a regular budget constraint to reach *indirect utility* function

$$v_{ij} = \max_{C_i, H_i} u_{ij} = F(A_i^r, A_j^w, w_j, p_i, q_i, d_{ij}) \cdot z_{ij}. \tag{2}$$

With optimal C_i and H_i, utility now depends on the unit price of consumption p_i, the unit price of residential floorspace q_i, and the wage (w_j) at workplace j. In RUM terminology, these prices are the attributes associated with the available alternatives in the location choice set. In the present case, the relevant attributes are derived from an underlying utility maximisation problem instead of being intuitively selected by the researcher.

The functional form of $F(\cdot)$ is determined by the specification of $f(C_i, H_i)$ in the direct utility function. The vast majority of QSMs adopt the usual Cobb-Douglas form with fixed

[7] We mentioned earlier that the location choice in QSMs shows similarity with previous models by Alex Anas and co-authors. Note that after taking logs, utility in (1) becomes equivalent to household utility in Anas and Liu (2007).

expenditure shares β and $1-\beta$. This assumption leads to a convenient indirect utility function proportional to $w_j/(p_i^{\beta} \cdot q_i^{1-\beta})$. In more advanced models with horizontally differentiated consumption goods, C_i may represent a sub-utility function of varieties with CES specification, in which case p_i is the price index. See Monte et al. (2018) for an early QSM example with this specification borrowed from NEG.

A multiplicative utility function combined with Fréchet-distributed idiosyncratic shock yields, after solving the discrete utility maximisation problem, the following choice probability for ij:

$$\pi_{ij} = \frac{[F(A_i^r, A_j^w, w_j, p_i, q_i, d_{ij})]^{\varepsilon}}{\sum_r \sum_s [F(A_r^r, A_s^w, w_s, p_r, q_r, d_{rs})]^{\varepsilon}}. \tag{3}$$

Just like in a logit model, this choice probability is a fraction where the numerator depends on the attributes of ij itself, while the denominator is the sum of the same quantity for all choice alternatives. But, as opposed to the logit case, these quantities are not exponentials but ε power functions of the locational characteristics. The full derivation of π_{ij} is available in the Appendix of Ahlfeldt et al. (2015) and the theoretical literature of discrete choice modelling.

Following the gravity literature, QSM papers often refer to the numerator and denominator of π_{ij} as *bilateral resistance* and *multilateral resistance*. Indeed, F in the bilateral term can be decomposed into a product of purely origin-dependent and destination-dependent functions and an impedance term, i.e.

$$F(\cdot) = F_o(A_i^r, p_i, q_i) \cdot F_d(A_j^w, w_j) \cdot F_{od}(d_{ij}), \tag{4}$$

which makes the gravity nature of this choice probability apparent. The role of multilateral resistance is to capture the relative attractiveness of alternative residence–workplace combinations. That is, π_{ij} might decrease simply because a competing location pair becomes more appealing. The fact that random utility location choice models provide a micro-foundation for commuting gravity equations has been known in the literature since the 1980s; see Anas (1983).

Similar to standard logit models, the representative household's expected utility has a closed-form expression,

$$E[v_{ij}] = \Gamma\left(\frac{\varepsilon - 1}{\varepsilon}\right)\left[\sum_r \sum_s [F(A_r^r, A_s^w, w_s, p_r, q_r, d_{rs})]^{\varepsilon}\right]^{1/\varepsilon} = E[v_{rs}] \quad \forall\, i, j, r, s, \tag{5}$$

where Γ is the Gamma function. Note that the term is square brackets is the denominator in choice probability (3). Thus this equation resembles the *logsum* formula in standard logit discrete choice models (Train 2009).

There is no difference in how expected utility must be interpreted either. When a larger pool of N households is assumed to follow the representative behaviour and only differ in the realisation of the idiosyncratic part of utility z_{ij}, the law of large numbers implies that $N \cdot \pi_{ij}$ households choose i and j. The level of utility is not identical among those who chose i–j as z_{ij} takes a different value for everyone. However, in a large enough sample $E[v_{ij}]$ is the same for each residence–workplace combination and $N \cdot E[v_{ij}]$ is an unbiased measure of aggregate utility. When a policy intervention alters the choice probabilities and thus the spatial distribution of residences and workplaces, the change in $N \cdot E[v_{ij}]$ quantifies the impact of the intervention on household welfare.

Note that expected utility as defined in (5) is not a money-metric measure of utility and thus it cannot be compared directly to investment costs arising in monetary terms, for example. Ideally, $E[v_{ij}]$ should be normalised by the marginal utility of money to transform it into the monetary dimension. However, the marginal utility of money differs between original-destination pairs in this model. One possible solution is to compute an expectation of the marginal utility of money. Another option is to apply the *equivalent variation* as a money-metric measure of welfare. That is, to compute the amount of extra income in the original scenario which delivers the same increase in utility that the policy intervention achieves. This measure is derived in Koster (2023), for example.

3.2 General equilibrium: integrating transport and the urban economy

Just like other SCGE models, QSMs combine households' location choice behaviour with production sector(s) and a floorspace construction market. The interaction between workers and these sectors of the economy governs the mechanisms behind most of the endogenous location characteristics introduced earlier, including consumption and housing prices at residential locations and the wage at each workplace.

Figure 1 depicts the schematic layout of a general equilibrium model. The three types of agents interact through the labour, goods and floorspace markets where (partial) equilibria are characterised by market clearing prices and quantities. These linkages are represented by dashed and continuous lines in the figure. Similar to households, firms in the manufacturing and construction sectors perceive market prices and local geographical fundamentals and determine the output they produce via an internal optimisation process, normally profit maximisation.[8] For the sake of brevity, in this section, we describe the underlying economic rationale verbally.

The production sector uses labour, commercial floorspace, and potential intermediate inputs to produce consumption goods and services. They take the market clearing factor prices and the price of the commodity(ies) as given and determine their factor demand (employment and floorspace use) and output levels. In general equilibrium, these quantities equal the amount of labour and floorspace supplied by workers and the construction sector, and consumption, under the prevailing market prices, in each location. The literature diverges in terms of the specific assumptions applied to the production sector.

1. In urban models, including Ahlfeldt et al. (2015), firms in a location are represented by a single production function. Locations across the urban area produce a uniform costlessly traded output at a normalised market price, under perfect competition. Local productivity depends on a local fundamental and access to economic mass via agglomeration. Higher productivity in dense areas leads to higher wages and higher employment.
2. In a regional or international context costless trade is less realistic. Monte et al. (2018) and others in the QSM literature thus applied an NEG-style production sector in which horizontally differentiated varieties of the consumption good are produced in each location. They assume increasing returns to scale (by introducing a fixed production cost) and monopolistic competition. Thus, the cost of moving goods around in space determines the consumption price index and the markup (spatial monopoly power)

[8] Certain activities in the public sector may follow a different objective function, for example social welfare maximisation. Adding the public sector as a fourth segment of the economy is a promising extension not addressed in the literature so far.

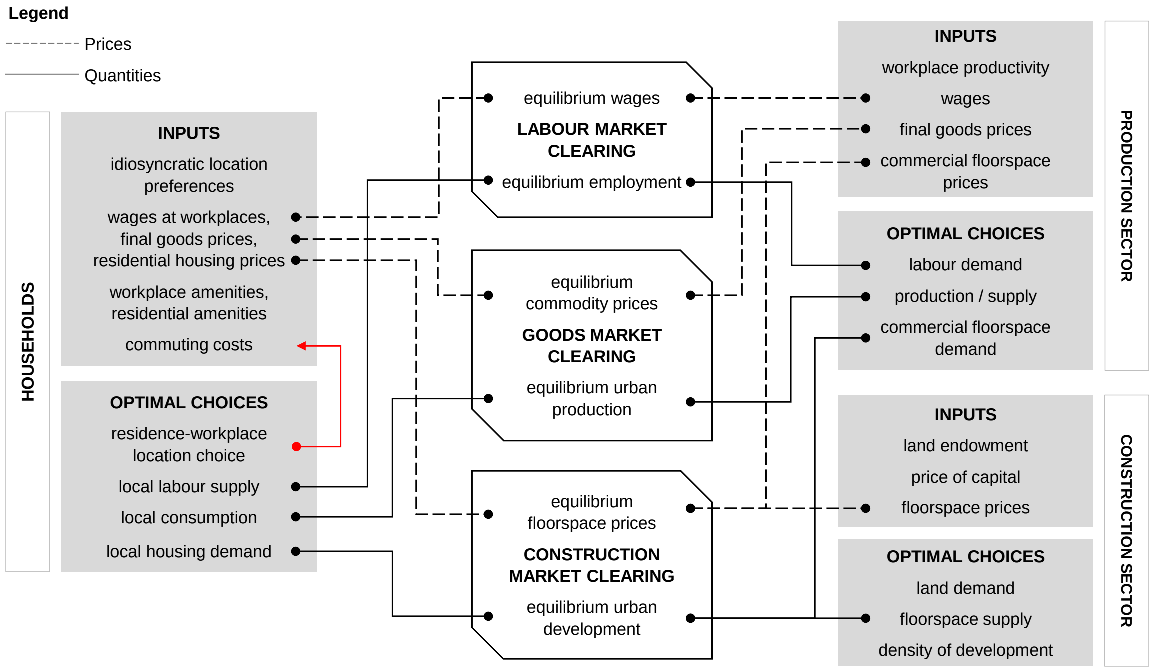

Figure 1: Schematic layout of an urban general equilibrium model

in each location, giving ground to agglomeration economies even in the absence of endogenous productivity.

3. One branch of regional QSMs extends the model to a two-tier production sector with input-output linkages between the producers of tradable intermediate varieties (i.e. tasks, business services) and the producers of non-tradable local goods/services. This approach enables the modeller to capture the impact of transport improvements on business-to-business connectivity. See, for example, Koster et al. (2023) in the context of high-speed rail in Japan.

The construction sector turns land and capital into floorspace, normally assuming a perfectly competitive profit-maximising market in each location. The outcome of the supply problem is the available floorspace and the density of development which may be constrained locally by regulation. By the density of development, we mean the ratio of floorspace supply and the developable land endowment in each location. Construction market clearing determines the equilibrium price and quantity of floorspace use. Who owns floorspace and what happens to the rent payments? This assumption may affect the welfare analysis significantly. Most QSMs assume for simplicity that landlords are immobile and spend all their rental income on consumption goods.

Many papers including Ahlfeldt et al. (2015) and Heblich et al. (2020) differentiate separate floorspace prices for residential and commercial uses, and allow their ratio to be a constant local characteristic recovered from data. They assume the wedge between the two prices is governed by local regulation, that is, differences in how residential and commercial floorspace uses are taxed.

Interestingly, the mainstream literature does not consider that transport infrastructure also consumes a part of the land endowment. This aspect has been considered in some of the earlier SCGE models including Anas and Liu (2007). Ducruet et al. (2023) are an exception in the QSM literature as they assume landlords can assign floorspace to transport use (i.e. port development) or regular residential use.

From the viewpoint of transport appraisal, the specification of the production and construction markets matters for two reasons. (i) Market failures in these sectors, agglomeration externalities and imperfect competition for example, may affect the welfare impact (also known as wider economic impact) of a transport policy. WEIs can be quantified more consistently, avoiding the threat of double counting, in a general equilibrium framework. (ii) Wages at the workplaces and consumption and housing prices at potential residential locations create micro-foundations for commuting behaviour. By explicitly modelling these prices we get more precise estimates of induced commuting demand. On top of that, to capture non-commuting traffic flows, for example business travel and freight flows, we need to consider the actions of the production side of the economy. The next section covers market failures first, while section 3.4 discusses interactions across space and traffic flows.

The schematic model depicted in Figure 1 can be extended by various additional aspects of household and firm behaviour. From the viewpoint of transport analysis, the micro-foundations of non-commuting trips might be especially relevant. For example, Fajgelbaum et al. (2023) assume that households' optimal choices also include the number of leisure trips to various destinations while business trips are a production factor and the time and monetary cost of travel are inputs for firms in the production sector as well. See further discussion on modelling transport in a QSM in Section 3.4 and the specificities of Fajgelbaum et al. (2023) in Section 5.

3.3 Market failures

The UK Transport Analysis Guidance[9] (TAG) distinguishes the following market failures deemed potentially relevant in the context of transport appraisal.

1. Product markets
   (a) Imperfect competition: lower monopoly mark-ups as competition intensifies with better connectivity.
   (b) Tax distortions: mismatch between gross and net prices, disincentivising economic activity.
   (c) Positive externalities from product variety.
2. Land markets
   (a) Land rationing: planning policies and land use/density restrictions imply divergence in the market price of developed and undeveloped land.
   (b) Imperfect competition in land markets.
   (c) Coordination failure: lower-than-optimal private investment in transport infrastructure after floorspace development.
3. Labour markets
   (a) Frictional unemployment: temporary unemployment due to frictions in job search.
   (b) Wage rigidities and structural unemployment: sticky prices in labour markets due to trade union bargaining power, wage regulation, uncertainty.

[9] Access online through http://www.gov.uk/guidance/transport-analysis-guidance-tag.

(c) Labour tax distorsions
(d) Monopsony buyers: monopsony power of large employers.

4. Agglomeration: Marshallian productivity externalities caused by connectivity contributing to effective economic density.

Current appraisal practice is restricted to an extended partial equilibrium approach. The economic model derives direct user benefits from a transport assignment model. Then, some of the wider economic impacts above – in particular, output chance in imperfectly competitive markets, labour supply impacts including move to more productive jobs, and agglomeration-generated productivity externalities – are calculated via simple approximation rules without modelling a new equilibrium in the product, land and labour markets following the intervention. This approach is simple enough for widespread practical application but it is restricted to market failures 1/a, 2/a, 3/c and 4. Moreover, the suspicion that some of these benefits emerging in non-transport markets may be captured in travellers' willingness to pay to commute cannot be unambiguously rejected.

As QSMs are spatial general equilibrium models, the double counting concerns unique to the extended partial equilibrium approach do not arise in their case. Also, the spatial general equilibrium approach is suitable to handle a wider set of market failures among the ones enlisted above in one integrated model. Some of the market failures are already incorporated in existing QSMs, and many others could be considered in reasonable extensions of the existing models.

3.3.1 Agglomeration economies

The externalities linked to agglomeration economies, both productivity and amenity spillovers, are well known and covered in the QSM literature. For instance, this has been the main subject of Ahlfeldt et al. (2015), the first published quantitative urban model. They followed the previous empirical literature on agglomeration (see Graham and Gibbons 2019, for a transport-oriented review) by assuming that agglomeration economies manifest through total factor productivity, that is, through a productivity shifter in the standard Cobb-Douglas production function:

$$y_j = A_j(\rho_j) H_j^{\alpha} L_j^{1-\alpha}, \tag{6}$$

where y_j is the output of workplace j, H_j and L_j are commercial floorspace and labour inputs, and α is the fixed expenditure share of floorspace use in production. Productivity shifter $A_j(\rho_j)$ is split into a *production fundamental* a_j representing local geographical characteristics and a function ρ_j, the measure of access to economic mass (ATEM).

$$A_j(\rho_j) = a_j \cdot \rho_j^{\lambda}, \tag{7}$$

where

$$\rho_j = \sum_k \exp(-\delta \tau_{jk}) \cdot L_k. \tag{8}$$

As elsewhere in the agglomeration literature, λ is the elasticity of productivity with respect to ATEM.[10] The ATEM measure is defined as a negative exponential function of travel cost τ_{jk} with decay parameter δ. Employment in location k, denoted by L_k, contributes less to ATEM

[10] Note that both λ and δ are generic parameters in the model, that is, they are not location-specific.

in j as travel time increases, while $\exp(-\delta\tau_{jk}) = 1$ when $j = k$ or the distance between j and k is very low.

Amenity spillovers, that is, agglomeration economies in the attractiveness of residence and workplace locations are incorporated in many QSMs. A_i^r and A_i^w, as defined in equation (1), can be decomposed identically into a fundamental element and a function of effective residential or employment densities following the expressions introduced above.

Quantifying the strength of agglomeration economies is an integral part of the QSM estimation process. In the first step, the unique value of A_j is recovered from observed data using the invertibility property of the model. See further details of the methodology in Section 4.2. Second, the agglomeration elasticity λ, the distance decay δ, and their equivalents for amenity externalities are estimated with other generic parameters in structural estimation exercises. As detailed in Section 4.3, this is performed either simultaneously using the Generalised Method of Moments or in a series of regressions exploiting the recursive structure of the model. Finally, when the vector of A_j's, λ and δ are known, the ATEMs can be computed for each location and we can solve for the unknown a_j local fundamentals by rearranging equation (7).

This quantitative approach assures that the econometric specification in which λ and δ are estimated is fully consistent with the general equilibrium model we use for welfare evaluation, and that they measure the strength of agglomeration forces in the local context. It is not surprising that the agglomeration elasticity is re-estimated at nearly every QSM application since Ahlfeldt et al. (2015), including the UK-oriented ones in Heblich et al. (2020), Dericks and Koster (2021), and Koster (2023).

3.3.2 Imperfect competition, land use and tax distortions

Market structure. Market failures 1/a and 1/c in the product markets are closely related to each other. In the standard NEG framework with monopolistic competition monopoly mark-ups as well as the number of product varieties are endogenous outcomes. Monopolistic competition with increasing returns and CES preferences (love of variety on the consumer side) is compatible with the QSM approach since Monte et al. (2018). Thus, the welfare measures derived in a QSM do capture the associated distortions, including output change in the imperfectly competitive product markets.

By contrast, land markets and the construction sector are normally considered perfectly competitive in existing QSMs.[11] This assumption is made to simplify the analytical work by relying on the standard Cobb-Douglas functional form. Assuming an imperfectly competitive construction sector to add market failure 2/b is not prohibited theoretically. But further analysis is needed to assess the quantitative relevance of this assumption to justify the move to a more complex specification. An even more complex option is to introduce separate construction and floorspace ownership markets, following the SCGE model of Anas and Liu (2007). The prevalence of model invertibility and the possibility of econometric estimation have not been proven in these setups so far.

Land markets. Land rationing policies (2/a) are an integral part of QSMs through local density limits, as explained on page 12. Indeed, land regulations are distortionary, and QSMs are suitable to measure this deadweight loss. Koster (2023) devotes his paper to this particular

[11] Some of the LUTI models are more "realistic" in this respect.

issue with an application to greenbelts in England. Tsivanidis (2023) also simulate the impact of a mass transit policy with and without density restrictions near the new layout. However, density limits are normally considered as an exogenous constraint in QSMs. It is questionable that a transport intervention can be the direct *cause* of the relaxation of land use regulation. Thus, we are not convinced that this market failure can lead to welfare gain in the assessment of a transport project.

We are not aware of any QSMs or even SCGE models in which the coordination failure 2/c above is modelled explicitly. More research is needed to assess its quantitative relevance and technical feasibility from a modelling point of view.

Labour markets. A transport intervention may affect the labour market through three channels: (i) job moves, "move to more productive jobs" in TAG terminology, (ii) labour force participation, and (iii) more subtle changes in individual labour supply. As workplace choice is endogenous in QSMs, the framework is suitable to capture channel (i) and quantify any changes in the spatial distribution of employment. Note, however, that QSMs only derive choice probabilities applied to a pool of ex-ante homogeneous workers. Thus, this framework is not suitable for 'tracking' individuals and identifying the source of an increase in employment in one specific location.

Fostering labour force participation and job creation might be an important aim of a transport scheme but this spatial outcome is rarely modelled in QSMs. One relevant stream of work is Delventhal et al. (2022) and Delventhal and Parkhomenko (2023) who consider endogenous working from home decisions in a QSM. Technically, unemployment can be modelled as an outside option in the workplace choice problem. We believe QSMs could provide very precise predictions on the job creation effect of transport policies. In channel (iii), one may be interested in the intensive margin of individual labour supply. In this case, existing QSMs remain restrictive by assuming that individual labour supply is inelastic, i.e. that each worker supplies the same number of workdays with the same working hours irrespective of the workplace and commuting distance. Hörcher and Graham (2023) formulate a QSM in which individual labour supply becomes endogenous via the leisure-labour trade-off well-known from Anas and Liu (2007) and other stylised-space models of urban economics (see e.g. DeSerpa 1971, Arnott 2007).

Note that the above-mentioned general equilibrium mechanisms in labour markets do not imply market failures per se. Appropriate modelling is only needed to get reliable estimates of labour supply and commuting demand in policy appraisal. Frictional and structural unemployment (3/a and 3/b) are difficult to distinguish in a QSM context. Monopsony power among potential employers (3/d) is not modelled in QSMs. Similar to imperfect competition in land markets, we believe it is possible to relax the perfect competition and price taking assumption in labour markets, but this extension requires more research.

Distortionary taxes. Surprisingly, value-added and labour taxes (market failures 1/b and 3/c) are normally not accounted for in mainstream QSMs. Only real estate taxes are considered indirectly through the distinction between residential and commercial floorspace use. It is surprising that the existing literature neglects other taxes as it seems particularly straightforward to impose a fixed tax rate on gross product prices and wages. The quantitative impact of such an extension is an interesting subject within reach for future research.

3.4 The transport market

Arguably, the way transport is modelled in the QSM literature is far from the state-of-the-art of transport science. We identify limitations in

(i) how transport cost interacts with other sources of utility,

(ii) the dimension of transport costs,

(iii) the range of trip purposes,

(iv) the available transport modes, and mode and route choice,

(v) the exogeneity of transport supply in most of the QSM applications.

As we noted in Section 3.1 already, the utility function frequently adopted in QSMs is based on the *iceberg* specification. That is, the disutility of commuting enters utility multiplicatively, implicitly assuming that this source of inconvenience is proportionate to the other determinants of utility such as consumption and floorspace availability. Whilst the iceberg assumption is common in the international trade literature and new economic geography where passenger transport is of limited relevance, its use is unconventional in urban and transport economics.

Limitation (ii) refers to the fact that most QSM studies do not distinguish the monetary and temporal dimensions of transport costs. In fact, some of the most influential papers such as Ahlfeldt et al. (2015) and Heblich et al. (2020) consider travel time only, ignoring monetary expenditures entirely. This assumption precludes the analysis of transport pricing policies. Even if the modeller attempted to replace time with a measure of generalised time in these models, a reliable estimate of the value of time would be needed to transform monetary expenses into generalised time.

Previous stylised-space general equilibrium transport models (e.g. Arnott 2007, Hörcher et al. 2020) as well as the SCGE model of Anas and Liu (2007) avoid the iceberg assumption and handle both temporal and monetary expenses by defining separate temporal and monetary budget constraints in the household problem. The ratio of the associated Lagrange multipliers (shadow prices of time and money) is an endogenous measure of the value of travel time. Hörcher and Graham (2023) apply this approach in the QSM framework, thus showing a feasible path to relaxing limitations (i) and (ii) above.

Item (iii) raises the obvious point that QSMs are almost entirely restricted to one trip purpose, commuting, while in today's travel patterns commuting accounts for far less than half of urban travel. This share has likely decreased further since the Covid pandemic due to remote work. There is room to incorporate leisure and business travel in spatial general equilibrium, but proving the prevalence of the advantageous QSM properties (whether the model remains invertible and suitable for econometric estimation) under such extensions requires further research. An initial attempt to model leisure and business travel in the context of long-distance rail transport and a QSM is documented in Fajgelbaum et al. (2023).

Existing QSMs do not intend to reproduce the transport network with the complexity of a full-blown traffic model. Certain models are restricted to one mode of transport (e.g. Fajgelbaum and Schaal 2020), others consider multiple but mode choice is restricted to the path with the shortest travel time (e.g. Allen and Arkolakis 2014, Donaldson 2018). Road congestion is ignored in every paper published in the 2010s. Allen and Arkolakis (2022) take a significant step by adding endogenous travel times and stochastic route choice to the QSM framework.

They also prove the conditions under which the new model leads to a unique equilibrium and maintains the invertibility property. Tsivanidis (2023) integrates a nested logit mode choice module into the Fréchet location choice framework, thus bringing QSMs closer to standard transport models.

Finally, in item (v) we note that transport supply is exogenous in the QSM literature, so the inherent rationale behind transport service provision, price setting and the choice of capacity are ignored in this tradition. This may be particularly relevant in liberalised transport markets such as the freight sector or long-distance passenger transport.

Can we simply plug a state-of-the-art traffic assignment model into the QSM framework to avoid the limitations above? That is possible in principle but it is likely that most of the advantageous analytical properties of QSMs would become intractable after such an integration. Our main concern is that the mainstream traffic assignment models are black boxes from an analytical point of view: their equilibrium properties, including the uniqueness of traffic outcomes for a given spatial equilibrium, cannot be proven analytically. Thus, we would lose a significant part of the benefits that make QSMs attractive among researchers.

As QSMs are general-purpose spatial models rather than ones built for direct use in transport analysis, the limitations above certainly cannot be treated as deficiencies. Our conclusion is that the QSM approach requires further work to be adjusted for the purpose of transport appraisal. It is likely that even after this work is completed, the resulting model will not reach the same level of detail in the theoretical model that the most advanced SCGE transport models, e.g. Anas and Chang (2023), do achieve. We believe, however, that a transport-oriented QSM might deliver a more ideal balance between model realism and adequate parameter estimation than any previous solutions in the literature.

3.5 Welfare as a model output

One of the undeniable benefits of the currently used partial equilibrium appraisal method is that the overall welfare effect of the investigated policy can be decomposed into additive benefit and cost elements. For example, one can decompose direct user benefits into travel time savings, vehicle operating cost savings, the benefit of induced trips, and other quantities. Such decomposition helps decision-makers understand why a particular transport improvement is beneficial for society and adds a sense of realism to the appraisal exercise which also improves our trust in the method.

Such additive decomposition is more difficult in the QSM framework. In the underlying discrete choice model, location choice decisions are governed by a large set of variables directly affected by the intervention or indirectly affected through the post-intervention reorganisation of the spatial economy. The relationship between these determinants is multiplicative and non-linear. Equation (5) is an appropriate measure of the net user benefit from a theoretical point of view, but we cannot decompose it into additive elements. That is, we cannot tell the extent to which higher wages, lower commodity prices and lower housing rents contribute to the benefit perceived by the representative household. Moreover, the model assumptions imply that in equilibrium, expected utility is the same across all OD-pairs. In the new equilibrium after the intervention, expected utility is again identical across the OD-pairs, but this common expected utility is different from the one before the policy. So everyone is affected by the policy but to the same extent. This implies that spatial equity analysis is meaningless in the baseline framework.

There are remedies to this problem, however. Benefits and costs can be quantified by fully restricting and then gradually relaxing the general equilibrium mechanisms of the model. It may be possible to reproduce the three levels of analysis outlined in TAG in a QSM. Level 1 analysis would mean that we keep the OD-matrix as well as the endogenous productivity and amenity variables unchanged relative to the initial spatial equilibrium and quantify expected utility for each residence–workplace combination individually. This step would also allow us to measure the spatial distribution of benefits in the short-run, before any relocation happens. On Level 2 one may recalculate ρ_i (the ATEM measure) and the endogenous productivity and amenity values for each location, and compute again expected utility for each OD pair. The difference between Level 2 and Level 1 results equals the contribution of static agglomeration benefits in addition to direct user benefits. Finally, Level 3 would involve a new spatial equilibrium with full relocation. This readjusts expected utility to the same level everywhere and allows the analyst to quantify the additional welfare effect of relocation on top of the static impacts.

The gradual approach outlined here reveals spatial distributional impacts. Note that equity can be measured in other dimensions as well. For example, Tsivanidis (2023) distinguishes income groups and Warnes (2021) has multiple skill groups. Having multiple types of representative households increases the data need of the modelling effort, however.

4 Quantifying the model

Calibrating large-scale general equilibrium models is a challenging task which, if poorly performed, may undermine our level of confidence in the resulting welfare measures and policy recommendations. This is especially the case for spatially explicit models that rely on hundreds of parameters capturing geographical features. QSMs avoid the need for ad-hoc parameter selection as their theoretical structure creates a one-to-one relationship between unknown geographical fundamentals and the observed data; this methodology is described in Section 4.2. Moreover, the remaining generic (non-geographical) parameters can be estimated in structural econometric exercises consistent with the equilibrium conditions of the model; see Section 4.3. Before we turn to these empirical methods, let us summarise the data need of the QSM approach.

4.1 Data

The increasing availability of large-scale spatial data sources is undoubtedly one of the main catalysts of the emergence and success of the QSM literature. Practice-ready QSMs split the geographical space into a discrete set of spatial units (districts/zones/blocks) covering a predefined area of developable land. Spatial units are often called *locations* in this literature. The centroids of the spatial units are connected by a graph representation of the transport network where each transport link has unique attributes, including capacities, and there may be transport nodes (intersections, hubs) besides the location centroids.

There is no known restriction on the resolution of the spatial model from a technical or computational point of view. An extreme example is Ahlfeldt et al. (2015) who handled 15,937 blocks and thus around 254 million origin-destination pairs independently on a conventional desktop. Subsequent QSM have not reached this level of spatial granularity. Evidently, models become less accurate at high resolution and thus this is a trade-off in the choice of the number of spatial units. In most cases, this choice is heavily influenced by the available data sources and spatial units often overlap with the relevant breakdown used by the data provider (e.g. output areas, census tracts, etc). Any mismatch between the statistical units used by different data providers can be handled by slicing and merging them using advanced GIS techniques.

Ideally, a seamlessly quantified QSM utilises disaggregate spatial data on (i) travel patterns, i.e., an origin-destination matrix of commuting volumes, (ii) congestion-free travel conditions and the capacity of road and public transport links, (iii) wages by residence or work-place, (iv) commercial and residential rents, and (v) the density of development and the share of land use purposes. Some of these data sources are publicly available in many countries (ii and v), others are provided by transport agencies and statistical offices on request (i, iii, iv). Beyond these traditional, mostly census-based data sources the QSM literature is gradually turning towards new automated data such as travel patterns retrieved from smartphone GPS data in Miyauchi et al. (2022).[12]

Since most of the leading QSM analyses are performed in a historical context, numerous techniques have been developed to circumvent the absence of certain data sources in the list above. The "exact hat algebra" approach developed by Dekle et al. (2007) enables the calculation of counterfactual equilibria based on the known relative change in a subset of observed variables, backing out other data sources that the analyst does not observe for the counterfactual time period (see e.g. Heblich et al. 2020). Recent research efforts show that QSMs can be successfully calibrated with very limited data sources as well. For example, Makovsky (2023) estimates agglomeration elasticities for the Czech Republic where disaggregate wage data are not available. He recovers the spatial distribution of wages from a census-based commuting matrix, using the QSM assumptions outlined below in Section 4.2. Sturm et al. (2023) recover commuting patterns from smartphone data and commuting times from a route planner, effectively claiming that a QSM can be calibrated anywhere in the developed as well as developing world.

4.2 Model inversion: recovering local geographical fundamentals

The full parameter set of a spatial model can be split into two groups. There may be several vectors of parameters quantifying the exogenous attributes of the discrete locations in the model. Stylised-space models often assume that the underlying geography is homogeneous and therefore there is no need to differentiate locations by geographical characteristics. This assumption is unrealistic when certain locations are clearly superior in terms of attractiveness for various activities. However, the quantitative measurement of first-nature geographical features in utility (or money) terms is particularly difficult and prone to omitting important determinants of such valuations.

QSM papers tackle this challenge by treating first-nature geographical features as *structural residuals* of the model. That is, they assume that the economic model captures a part of the determinants of households' and firms' location choices and what the model is unable to explain

[12] The move to big data is also a prevalent theme in the transport modelling and planning sector.

must be attributed to geographical advantages and disadvantages.[13] This assumption can only be operationalised if there is a one-to-one relationship between the vector local fundamentals and the observed data, i.e. if only one vector of parameters can turn the observed data into an equilibrium outcome of the model. The literature calls this process model inversion. The proofs of invertibility vary by model specification.

The general idea behind the process of inversion is that the model's equilibrium conditions can be ordered in such a way that in each step, only one unknown vector of local fundamentals is expressed in function of observed data.

Let us illustrate this idea on the recovery of A_j^w, the attractiveness of model locations as workplaces. It can be shown that the condition probability that the representative worker who lives in i commutes to j is

$$\pi_{ij|i} = \frac{[F_d(A_j^w, w_j) \cdot F_{od}(d_{ij})]^\varepsilon}{\sum_s [F_d(A_s^w, w_s) \cdot F_{od}(d_{is})]^\varepsilon}, \tag{9}$$

in which we see that the attributes of the origin, $F_o(A_o^r, p_i, q_i)$ in equation 4, are cancelled out but the rest of the expression remains similar to π_{ij} in (3). When applied to a large pool of households, this conditional probability creates a link between workplace employment in j and the residential population of i, denoted by N_j^w and N_i^r, respectively.

$$N_j^w = \sum_i \pi_{ij|i} \cdot N_i^r \tag{10}$$

Now, after substituting expression (9) above into (10), it is clear $F_d(A_j^w, w_j)$ is identical for the origins i and therefore it can be taken out of the summations:

$$N_j^w = [F_d(A_j^w, w_j)]^\varepsilon \sum_i \frac{[F_{od}(d_{ij})]^\varepsilon}{\sum_s [F_d(A_s^w, w_s) \cdot F_{od}(d_{is})]^\varepsilon} N_i^r. \tag{11}$$

By rearranging this equation, that is,

$$F_d(A_j^w, w_j) = \left\{ \frac{1}{N_j^w} \sum_i \frac{[F_{od}(d_{ij})]^\varepsilon}{\sum_s [F_d(A_s^w, w_s) \cdot F_{od}(d_{is})]^\varepsilon} N_i^r \right\}^{-\frac{1}{\varepsilon}}, \tag{12}$$

we express $F_d(A_j^w, w_j)$ in function of vectors N_j^w and N_i^r, and the commuting cost matrix d_{ij}. Ahlfeldt et al. (2015) shows that (12) has a unique solution so that only one realisation of vector $F_d(A_j^w, w_j)$ satisfies equation (12) for each j. They also show an iterative algorithm, e.g. the method of successive averages, converges to that solution. If $F_d(A_j^w, w_j)$ is simply the product of the local amenity and the wage, the vector of local amenities can be recovered in one last step: $A_j^w = F_d(A_j^w, w_j)/w_j$.

N_i^r and N_j^w can be computed as the row and column sums of a standard commuting matrix, the d_{ij} transport cost matrix is a regular input of transport modelling, and wages are normally also available for the selected geographical area. This implies that we have a one-to-one mapping between these observed data and a unique A_j^w value for each workplace in the model.

A similar process based on the conditional *residence* choice probability $\pi_{ij|j}$ allows us to recover A_i^r, the amenity level of residential locations. The vector of firm productivity A_j, introduced

[13] The use of structural residuals is common outside the spatial economics literature; see, for example, the "Solow residual" in the development economics literature.

earlier in Section 3.3.1, is expressed from the wage equilibrium condition of the production sector. These steps are repeated in nearly every QSM application. There are other local fundamentals defined in various models which are quantified using tailor-made methods that we do not review here.

Is this methodology entirely new? In the case of residence and workplace attractiveness variables, we found an earlier clue in the literature. Anas (1983) shows that the maximum likelihood estimates of alternative specific constants of a logit location choice models can also be recovered from aggregate data, using a commuting matrix without data on individual choice situations. Behrens and Murata (2021) show this link referring back to the original contributions made by McFadden (1974).

We close this section with a word of caution: treating local geographical fundamentals as structural residuals of the model offers a convenient solution to back-out missing information on what drives people's preferences for certain locations. In fact, this approach never fails as the calibrated local fundamentals ensure that the equilibrium conditions are always exactly satisfied. This overly advantageous property may hide errors stemming from model misspecification and noise in the data. In other words, the residuals may capture other important omitted effects unrelated to geographical advantage.

4.3 Structural estimation of generic parameters

Standard QSMs operate with a limited number of generic parameters, namely

(i) the Fréchet shaper parameter (ε) capturing the degree of heterogeneity in preferences,

(ii) a commuting cost semi-elasticity κ when commuting disutility in equation (1) is specified as $d_{ij} = \exp(\kappa\tau_{ij})$,

(iii) expenditure shares in the Cobb-Douglas utility and production functions,

(iv) the elasticity of substitution when utility (or one of the production functions) is CES,

(v) the agglomeration elasticity and distance decay parameters (λ and δ in Section 3.3.1), separately defined for productivity and amenity externalities.

Among these parameters, the Cobb-Douglas expenditure shares are normally borrowed from other empirical papers in the literature.[14] Furthermore, agglomeration economies are more relevant in an urban context while imperfect competition is more of a regional feature, so agglomeration and substitution elasticities are rarely estimated in one paper. This leaves us with around six to ten generic parameters to be estimated in one QSM application.

The original quantitative urban model of Ahlfeldt et al. (2015) did a very rigorous job by presenting two structural estimation approaches. In their more comprehensive approach, they estimated ε, κ, λ and δ and their amenity equivalents in one step, using the Generalised Method of Moments (GMM). The estimated parameter vector minimises the GMM objective function (moment function) derived from at least as many moment conditions for identification

[14] This practice may be criticised as it prevents QSMs from being entirely self-contained in terms of parameter estimation. However, expenditure shares are easy to conceptualise and observe directly in data, and the estimates in the literature are among the most robust empirical findings in applied economics. In our view, the restrictive Cobb-Douglas functional form is a more serious weakness of the QSM approach than its imperfect calibration.

as unknown parameters the model has. Moment conditions may be (i) market clearing equations of the model, (ii) orthogonality conditions in case of instrumental variables, or (iii) other moments describing the distribution of observed data. Subsequent quantitative spatial economic studies tackled endogeneity with different identification strategies depending on data availability. The instrumental variables approach is frequently used when a large-scale natural experiment such as the division and reunification of Berlin (Ahlfeldt et al. 2015) or random bombing destruction in London (Dericks and Koster 2021) is not available. Regular instruments of travel time in gravity equations include Euclidean distance, planned route, and least cost route instruments, while the estimation of agglomeration elasticities and distance decay follow a well-known literature reviewed by Graham and Gibbons (2019).

In a more restrictive approach Ahlfeldt et al. (2015) assumed that local productivities and amenities are fixed (no agglomeration economies). They estimated ε and κ by transforming the choice probability in equation (3), using the decomposition in (4), into the following log-likelihood function.

$$\ln \mathcal{L} = \sum_{ij} N_{ij} \ln \pi_{ij} = \sum_{ij} N_{ij} \left[K + FE_i^o + FE_j^d + \varepsilon \ln F_{ij}^{od} \right] \tag{13}$$

In this specification, K is the constant denominator of the location choice probability which is identical for each residence–workplace combination. FE_i^o and FE_j^d are fixed effects that capture all the origin and destination-specific determinants of the choice probabilities. The remaining F_{ij}^{od} is unique to i and j and normally restricted to the commuting disutility. Ahlfeldt et al. (2015) rely on the most established methods of gravity estimation in the international trade literature (in particular, on Santos Silva and Tenreyro 2006, Head and Mayer 2014) and estimate (13) via the Poisson and Gamma Pseudo Maximum Likelihood methods. Ahlfeldt et al. (2015) show in counterfactual simulations that the previous GMM approach, which considers endogenous agglomeration externalities, performs much better in explaining the spatial reorganisation of economic activity after the construction and demolition of the Berlin Wall. However, when only cross-sectional data are available and/or it is difficult to establish a complete identification strategy, the QSM literature often relies on reduced-form PPML and GPML gravity estimation techniques as well.

Is there a way to standardise the estimation of structural parameters in a general appraisal guidance? The most influential QSMs published in top economics journals rely heavily on unique data, often linked to historical events that induced an exogenous shock in the spatial organisation of economic activity. One cannot guarantee the availability of such data sources in any geographical context. Future research should explore credible identification strategies that do not necessitate unique (historical) data sources but ensure the unbiased estimation of the structural parameters amid endogeneity concerns.

4.4 Validation

Validation is key to establishing trust in a quantitative tool meant to support policy evaluation. A reliable spatial economic model applied in transport appraisal should predict the future reorganisation of economic activity in response to major policy interventions precisely and provide a correct estimate of the underlying welfare effects. To the best of our knowledge, such large-scale validation experiments, with an ex-ante prediction benchmarked against the ex-post causal impact of the policy, have not taken place in the QSM literature so far—especially not in a transport appraisal context. We believe this task on the research agenda is important but also time-consuming, as the life cycle of the full validation experiment of a transformational scheme may take more than a decade.

However, most of the previous papers in the QSM literature cover long time periods in a historical context. This evidence base shows that the methodology performs well in reproducing the spatial reorganisation of economic activity following the construction and fall of the Berlin Wall (Ahlfeldt et al. 2015), the expansion of the railway network in India (Donaldson 2018), and the evolution of Greater London after the steam engine revolution (Heblich et al. 2020) or World War II (Dericks and Koster 2021).

4.5 Uncertainty quantification

Since QSMs draw on a relatively small number of structural parameters, for which uncertainty can presumably be represented (e.g. via the parameter standard errors), it should be possible to develop a coherent framework for uncertainty quantification. For instance, it may be possible to compute confidence (or credible) intervals for key model outputs.

This would represent a useful development in appraisal modelling. All too often the general equilibrium approaches used to supplement appraisal have failed to provide an evaluation of model uncertainty. Rather, model results are typically presented deterministically, that is, in the absence of a rigorous attempt to quantify uncertainty in a coherent fashion.

Beside the use of the standard errors of internally estimated parameters, the uncertainty of key appraisal outputs such as the NPV and BCR can be quantified via randomised numerical (also known as Monte Carlo) simulations. In each step of this numerical approach a random value of the core input parameters is drawn from a predefined probability distribution. Thus, the method yields a probability distribution of the appraisal outputs as well, from which we can compute the probability that the project's performance exceeds critical thresholds.

5 Transport-related applications

The QSM papers published in leading economics journals complemented model development and estimation with a compelling historical context, a series of reduced-form econometric exercises to motivate the spatial model, and large-scale illustrative counterfactuals. Even though a wide range of place-based policies can be tested in this framework, transport-related applications are surprisingly popular in the literature. The removal of the entire rail network of London in Heblich et al. (2020) or the Shinkansen network in Japan in Koster et al. (2023) are shown to induce a remarkable welfare loss and the reorganisation of economic activity at a groundbreaking scale. Note, however, that these counterfactual scenarios are not reasonable policy options in present-day transport politics, so QSM authors have not aimed so far to mobilise their framework to assess pressing policy dilemmas in transport. In this section, we present a brief overview of the transport-related aspects of published and unpublished QSM studies.

Table 1 enlists six published peer-reviewed articles and six working papers with high visibility in the spatial economics community. The table compares these studies based on their geography, time horizon, trip purpose and transport mode availability for both passenger and freight (trade) movements, and the transport-related counterfactual scenario they simulate.

Our list includes seven urban models for Berlin, London, Seattle, Bogota, Buenos Aires, Tokyo and Mexico City. These papers focus mostly on passenger transport. Tsivanidis (2023) is often referred to as the first model in which stochastic mode choice is adequately represented and nested into the location choice problem. His analysis focuses on the welfare and displacement effect of Bus Rapid Transit (BRT) development in Bogota. In counterfactual simulations, he tests the efficiency of alternative BRT networks. Furthermore, he simulates the welfare effect of the original intervention assuming a relaxation of zoning restrictions around new stations, showing substantial net benefits associated with the combined policy.

Commuting is the dominant trip purpose in most of the urban models. Miyauchi et al. (2022) is a unique work in the sense that they introduce non-commuting trips as well, and their location choice model enables trip chaining, i.e. the realisation of an arbitrary sequence of trips linking the household's residential and workplace locations with other activity scenes. They show in a counterfactual simulation that allowing for trip chaining besides regular commuting has a substantial impact on welfare outcomes. Zárate (2022) is another recent example in which consumption (shopping) trips complement commuting.

Although Warnes (2021) limits the analysis to commuting, his model is remarkable as the only *dynamic* spatial equilibrium model in this list. While all the other papers perform comparative statics between the two equilibrium states before and after a policy intervention, Warnes (2021) builds and quantifies a model in which the transition between two equilibria is explicitly captured. That is, households' relocation decisions are assumed to take time and thus he is able to predict the trajectory through which the urban economy readjusts to a new equilibrium. This feature seems particularly appealing in a transport appraisal context where such transitions are normally ignored. Intuition suggests that the pace at which future benefits are realised may have a crucial impact on the economic case of transformative projects. Also, dynamic quantitative spatial modelling opens up the opportunity to model foresight in location decisions explicitly. In a transport context, this would explain why property prices often increase near transport improvements way before the new infrastructure is completed.

The model by Delventhal and Parkhomenko (2023) is regional in scale as it covers the US entirely but its structure resembles urban models, focusing on passenger commuting.[15] This study responds to several recent policy questions about remote work in the post-pandemic era. They make the working from home decision endogenous in the model, thus making commuting patterns dependent on the productivity and wage differences between remote and on-site work. In counterfactual simulations, they test the impact of various productivity gaps on urban spatial structure.

The remaining papers with a regional scope are quantitative economic geography models. Allen and Arkolakis (2014) and Donaldson (2018) consider multiple modes of freight transport but mode choice remains deterministic on the basis of shortest travel time (differentiated by mode-specific time valuations in Donaldson 2018). Ducruet et al. (2023) is a unique paper as their maritime model has a global coverage. In counterfactuals, they test the impact of containerisation on global wellbeing, and a specific policy focusing on port development.

Multiple papers investigate the aggregate efficiency of national infrastructure networks and

[15] This setup is similar to Monte et al. (2018), and influential study with no particular relevance to transport.

Table 1: Literature overview: peer-reviewed publications (above) and working papers (below)

Article	Geography	Time horizon	Passenger trips: purpose and mode	Trade; freight transport modes	Transport counterfactual
Allen and Arkolakis (2014)	Regional: U.S.	Static: early 2000s	n.a.	Road, rail, inland water	No Interstate Highways
Ahlfeldt et al. (2015)	Urban: Berlin	1936–2006	Commute: underground, suburban rail, tram, bus, car	n.a.	No car use in 2006
Donaldson (2018)	Regional: India	1870–1930	n.a.	Road, rail, inland and coastal shipping	n.a.
Fajgelbaum and Schaal (2020)	Regional: 24 European countries	Static: mid-2000s	n.a.	Road	Road expansion and reallocation
Heblich et al. (2020)	Urban: London	1801–1921	Commute: overground, underground, omnibus, tram	n.a.	No railway network
Allen and Arkolakis (2022)	Regional: U.S., Urban: Seattle	Static: early 2010s	Commute: road	Road	Welfare effect of link-level improvements
Tsivanidis (2023)	Urban: Bogota	1993–2018	Commute: BRT, bus, walk, car	n.a.	Relaxed zoning restrictions. Alternative network
Warnes (2021)	Urban: Buenos Aires	2011–2017	Commute: BRT, public transport, walk	n.a.	Alternative BRT network configurations
Miyauchi et al. (2022)	Urban: Tokyo	Static: 2019	Commute and 'non-commute': walk, bus, metro, railway, car	n.a.	No urban rail lines. No trip chaining.
Zárate (2022)	Urban: Mexico City	Static: various data sources 2000–2017	Commute and consumption trips: BRT, bus, metro, car, walking	n.a.	Non-transport policies to reduce informality
Delventhal and Parkhomenko (2023)	Regional and urban: United States	2019-2022	Commute: mode not specified	n.a.	Varying productivity for remote work
Ducruet et al. (2023)	Regional: Global	Static: 1990	n.a.	Maritime and abstract inland transport	No container technology; Maritime Silk Road policy
Koster et al. (2023)	Regional: Japan	1957–2015	Commute: urban rail, road; business: HSR, urban rail, road	n.a.	No HSR; no highways; full HSR network realised
Fajgelbaum et al. (2023)	Regional: California, US	Static: latest data prior to 2020	Commute, leisure, business. Air, HSR, car, public transport, bicycle/walk	n.a.	California HSR line completion; alternative optimal alignments

local capacity expansions. Fajgelbaum and Schaal (2020) combine the QSM approach with an optimisation algorithm through which they reoptimise national road networks in 24 European countries taking general equilibrium implications through trade into account. Allen and Arkolakis (2022) compute the welfare elasticity of link-level capacity expansion in the US highway network to identify the bottlenecks where adding more lanes generates the greatest net benefit for society. Congestion delay is endogenous in both Fajgelbaum and Schaal (2020) and Allen and Arkolakis (2022).

Fewer studies focus on railways and other passenger transport modes in a regional context. Koster et al. (2023) broke this trend by modelling high-speed rail in Japan. As explained in Section 3, this paper considers both commuting and business travel, where demand for the latter is derived from firm-to-firm interactions in a two-tier supply chain. In a counterfactual simulation they show that the removal of the Shinkansen network would cause more harm to the Japanese economy than the removal of the entire highway network. Fajgelbaum et al. (2023) is another QSM application focusing on high-speed rail in California, US. Their microfoundation for business travel is somewhat simpler as they consider the number of business trips as a factor in firms' production function. Whilst it is convincing that firms become more productive due to the opportunity to travel, we believe this specification is more limited than the one by Koster et al. as it fails to capture the web of supply chain relationships and the way business travel densifies it. Not, however, that Fajgelbaum et al. (2023) delivers a series of other insights relevant to transport policy, including the first truly spatially explicit model of the political economy of infrastructure development.

We close this section by concluding that transport policies are frequent subjects in the QSM literature but their purpose is usually illustrative. Allen and Arkolakis (2022) sets the goal of analysing the welfare effects of transport policies explicitly, but their approach is still very far from the contemporary appraisal practice. Part of the reasons is the lack of cross-fertilisation between the spatial and transport economics communities, i.e. QSM authors are not familiar with the mainstream transport literature. In contrast, transport researchers have not fully recognised QSMs as a transport modelling framework yet.

6 Evaluation and conclusions

The analysis in this review project is summarised in the following concluding statements.

- The QSM methodology is a promising tool for transport appraisal, that is, the ex-ante welfare analysis of transport interventions. However, existing QSM studies have not set this aim so far as they address more general spatial economic research questions. The main limitations hindering their direct adaptation in transport appraisal are related to the simplistic way the transport sector is modelled in these studies. In particular, the *iceberg* specification of travel disutility is not consistent with transport economic theory, non-commuting trips are mostly ignored, and route and mode choice under congestion is inadequately modelled.
- QSMs might be especially useful to better understand the general equilibrium impact of transformative schemes. For small interventions, the partial and general equilibrium welfare effects are likely very similar. However, one may envisage a model framework in which Level 1, 2 and 3 impacts of a scheme could be modelled simultaneously by gradually activating the general equilibrium features in the coding environment. This

way one could avoid the need for two distinct model frameworks and to decide a priori which interventions must be analysed in a QSM.

- In this review we identified a trade-off between the richness of model features, the economic underpinning of the model, and the robustness of the empirical evidence behind parameter calibration/estimation. In this spectrum LUTI models are the richest in spatial features at the expense of both theoretical and empirical consistency, SCGE models are theoretically robust but difficult to calibrate, while QSMs strike a healthy balance between detailing and reliable parameter estimation, enabling the use of causal econometric tools.
- QSMs have not been deployed so far to form ex-ante policy recommendations in transport. However, a diverse set of existing ex-post applications reviewed in Section 5 provide munition for future deployment in appraisal.
- The QSM framework captures direct user benefits and the wider economic impacts currently considered in TAG in a theoretically integrated fashion, thus neutralising the double counting concerns that cannot be fully rejected in an extended partial equilibrium model. It is possible to derive Level 1, 2 and 3 performance metrics from one model. However, the general equilibrium framework may not be suitable to additively separate benefit items to the extent we currently do in TAG. For this reason, a QSM will provide alternative complementary evidence to partial equilibrium appraisal, but will not replace it.
- More specifically, existing QSMs feature endogenous productivities and amenities governed by agglomeration and an NEG-style monopolistic competition mechanism in the production sector(s) of the spatial model. Furthermore, adding various tax wedges and land use externalities is easy to implement in this framework. This creates reliable microfoundations to calculate the wider economic impacts conceptualised in TAG.
- To the best of our knowledge, no country has been using QSMs for transport appraisal so far.
- The quantitative general equilibrium approach significantly reduces the uncertainty of parameter calibration by allowing the direct estimation of the most relevant parameters from data. Beyond that, the general equilibrium approach does not provide a solution to handle other sources of uncertainty, specifically. We believe that numerical sensitivity testing and scenario analysis would remain indispensable in a QSM-based transport appraisal as well.
- The imposition of rigorous procedures for model validation and uncertainty quantification will be necessary to establish QSMs as a dependable tool for transport appraisal.

Quantitative spatial economics is a nascent concept, especially when it comes to transport economic applications. In other words, the method is not practice-ready yet and it must undergo academic development in the coming years. Government could play a pivotal role by supporting this evolution by sharing data, practical experience and financial resources with the research community involved in transport-oriented QSM development.

References

Acheampong, R. A. and E. A. Silva (2015). Land use–transport interaction modeling: A review of the literature and future research directions. *Journal of Transport and Land use 8*(3), 11–38.

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

Allen, T. and C. Arkolakis (2014). Trade and the topography of the spatial economy. *The Quarterly Journal of Economics 129*(3), 1085–1140.

Allen, T. and C. Arkolakis (2022). The welfare effects of transportation infrastructure improvements. *The Review of Economic Studies 89*(6), 2911–2957.

Allen, T., C. Arkolakis, and X. Li (2020). On the equilibrium properties of network models with heterogeneous agents. Technical report, National Bureau of Economic Research.

Anas, A. (1983). Discrete choice theory, information theory and the multinomial logit and gravity models. *Transportation Research Part B: Methodological 17*(1), 13–23.

Anas, A. (2020). The cost of congestion and the benefits of congestion pricing: A general equilibrium analysis. *Transportation Research Part B: Methodological 136*, 110–137.

Anas, A. and H. Chang (2023). Productivity benefits of urban transportation megaprojects: A general equilibrium analysis of «Grand Paris Express». *Transportation Research Part B: Methodological 174*, 102746.

Anas, A. and I. Kim (1996). General equilibrium models of polycentric urban land use with endogenous congestion and job agglomeration. *Journal of Urban Economics 40*(2), 232–256.

Anas, A. and Y. Liu (2007). A regional economy, land use, and transportation model (RELU-TRAN©): formulation, algorithm design, and testing. *Journal of Regional Science 47*(3), 415–455.

Anas, A. and R. Xu (1999). Congestion, land use, and job dispersion: a general equilibrium model. *Journal of Urban Economics 45*(3), 451–473.

Arnott, R. (2007). Congestion tolling with agglomeration externalities. *Journal of Urban Economics 2*(62), 187–203.

Behrens, K. and Y. Murata (2021). On quantitative spatial economic models. *Journal of Urban Economics 123*, 103348.

Dekle, R., J. Eaton, and S. Kortum (2007). Unbalanced trade. *American Economic Review 97*(2), 351–355.

Delventhal, M. and A. Parkhomenko (2023). Spatial implications of telecommuting. Available at SSRN 3746555.

Delventhal, M. J., E. Kwon, and A. Parkhomenko (2022). Jue insight: How do cities change when we work from home? *Journal of Urban Economics 127*, 103331.

Department for Transport (2005). Transport, wider economic benefits, and impacts on gdp. Discussion paper, July 2005, UK Department for Transport.

Dericks, G. H. and H. R. Koster (2021). The billion pound drop: the Blitz and agglomeration economies in London. *Journal of Economic Geography 21*(6), 869–897.

DeSerpa, A. C. (1971). A theory of the economics of time. *The Economic Journal 81*(324), 828–846.

Donaldson, D. (2018). Railroads of the Raj: Estimating the impact of transportation infrastructure. *American Economic Review 108*(4-5), 899–934.

Ducruet, C., R. Juhász, D. K. Nagy, and C. Steinwender (2023). All aboard: The effects of port development. NBER Working Paper Number 28148.

Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica 70*(5), 1741–1779.

Fajgelbaum, P. D., C. Gaubert, N. Gorton, E. Morales, and E. Schaal (2023). Political preferences and the spatial distribution of infrastructure: Evidence from California's high-speed rail. NBER Working Paper Number 31438.

Fajgelbaum, P. D. and E. Schaal (2020). Optimal transport networks in spatial equilibrium. *Econometrica 88*(4), 1411–1452.

Fosgerau, M. and M. Bierlaire (2009). Discrete choice models with multiplicative error terms. *Transportation Research Part B: Methodological 43*(5), 494–505.

Fujita, M. and P. Krugman (2004). The new economic geography: Past, present and the future. In *Fifty years of regional science*, pp. 139–164. Springer.

Fujita, M., P. R. Krugman, and A. Venables (1999). *The spatial economy: Cities, regions, and international trade.* MIT press.

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

Head, K. and T. Mayer (2014). Gravity equations: Workhorse, toolkit, and cookbook. In *Handbook of International Economics*, Volume 4, pp. 131–195. Elsevier.

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

Helpman, E. (1998). The size of regions. In D. Pines, E. Sadka, and I. Zilcha (Eds.), *Topics in Public Economics: Theoretical and Applied Analysis*, pp. 33–54. Cambridge: Cambridge University Press.

Hörcher, D., B. De Borger, W. Seifu, and D. J. Graham (2020). Public transport provision under agglomeration economies. *Regional Science and Urban Economics 81*, 103503.

Hörcher, D. and D. J. Graham (2023). Endogenous commuting time valuations and the leisure-labour trade-off in quantitative spatial modelling. Working Paper. Imperial College London.

Hunt, J. D. and D. C. Simmonds (1993). Theory and application of an integrated land-use and transport modelling framework. *Environment and Planning B: Planning and Design 20*(2), 221–244.

Koster, H. R. A. (2023). The welfare effects of greenbelt policy: Evidence from England. Forthcoming in The Economic Journal.

Koster, H. R. A., K. Hayakawa, T. Tabuchi, and J.-F. Thisse (2023). High-speed rail and the spatial distribution of economic activity: Evidence from Japan's Shinkansen. Working Paper. Vrije Universiteit Amsterdam.

Krugman, P. (1991). Increasing returns and economic geography. *Journal of political economy 99*(3), 483–499.

Lowry, I. S. (1964). A model of metropolis. Technical report, Rand Corporation Santa Monica.

Lucas, R. E. and E. Rossi-Hansberg (2002). On the internal structure of cities. *Econometrica 70*(4), 1445–1476.

Makovsky, L. (2023). Estimating agglomeration economies without wages: A quantitative spatial model approach exploiting commuting patterns. Working paper, The London School of Economics.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers in Econometrics*, pp. 105–142. Academic Press, New York.

McFadden, D. (1978). Modelling the choice of residential location. In A. Karlqvist, L. Lundqvist, F. Snickars, and J. Weibull (Eds.), *Spatial Interaction Theory and Residential Location*, pp. 75–96. North Holland, Amsterdam.

Miyauchi, Y., K. Nakajima, and S. J. Redding (2022). The economics of spatial mobility: Theory and evidence using smartphone data. NBER Working Paper Number 28497.

Monte, F., S. J. Redding, and E. Rossi-Hansberg (2018). Commuting, migration, and local employment elasticities. *American Economic Review 108*(12), 3855–90.

Nagy, D. K. (2022). Quantitative economic geography meets history: Questions, answers and challenges. *Regional Science and Urban Economics 94*, 103675.

Paulley, N. J. and F. V. Webster (1991). Overview of an international study to compare models and evaluate land-use and transport policies. *Transport Reviews 11*(3), 197–222.

Proost, S. and J.-F. Thisse (2019). What can be learned from spatial economics? *Journal of Economic Literature 57*(3), 575–643.

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

Robson, E. N., K. P. Wijayaratna, and V. V. Dixit (2018). A review of computable general equilibrium models for transport and their applications in appraisal. *Transportation Research Part A: Policy and Practice 116*, 31–53.

Santos Silva, J. M. C. and S. Tenreyro (2006). The log of gravity. *The Review of Economics and statistics 88*(4), 641–658.

Sturm, D., K. Takeda, and A. Venables (2023). How useful are quantitative urban models for cities in developing countries? Evidence from Dhaka. Working paper, The London School of Economics.

Tavasszy, L., M. Thissen, and J. Oosterhaven (2011). Challenges in the application of spatial computable general equilibrium models for transport appraisal. *Research in Transportation Economics 31*(1), 12–18.

Train, K. E. (2009). *Discrete choice methods with simulation.* Cambridge University Press.

Tsivanidis, N. (2023). Evaluating the impact of urban transit infrastructure: Evidence from Bogota's Transmilenio.

Warnes, P. E. (2021). Transport infrastructure improvements and spatial sorting: Evidence from buenos aires. Working Paper, Columbia University.

Wegener, M. (2011). Transport in spatial models of economic development. In A. de Palma et al. (Eds.), *A Handbook on Transport Economics*, pp. 46–66. Edward Elgar Cheltenham, Glos.

Zárate, R. D. (2022). Spatial misallocation, informality, and transit improvements: Evidence from Mexico City. Policy Research Working Paper, The World Bank, No. 9990.

Contents