Journal of Public Economics 259 (2026) 105683

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/pubec

A quantitative urban model for transport appraisal☆

Daniel Hörcher[a,b,*], Daniel J. Graham[b]

[a] Corvinus Institute for Advanced Studies, Corvinus University of Budapest, Budapest, Hungary
[b] Transport Strategy Centre, Department of Civil and Environmental Engineering, Imperial College London, London, United Kingdom

HIGHLIGHTS

- We develop a quantitative spatial model (QSM) for transport cost-benefit analysis.
- QSM welfare changes are decomposed into direct benefits and wider impacts.
- Case study: Greater London with 983 spatial units and the Elizabeth Line.
- CBA in spatial equilibrium reveals local impacts on labour and housing markets.
- We quantify standard errors for CBA results and various spatial outcomes.

ARTICLE INFO

Keywords:
Transport infrastructure
Cost-benefit analysis
Quantitative spatial economics
Agglomeration economies
Travel time valuation

ABSTRACT

Transport cost-benefit analysis (CBA) has evolved since its inception to become one of the most influential and ubiquitous applications of microeconomic theory, shaping billions of dollars of investment in the infrastructure sector. A key limitation of transport appraisal practice is its reliance on partial equilibrium (PE) models, which restricts the ability to quantify potentially transformative impacts outside the transport market. In this paper, we build on the principles of quantitative spatial economics and design an appraisal method that integrates the causal estimation of key parameters into an invertible spatial general equilibrium (SGE) model. Our specification yields travel time valuations that are micro-founded through an explicit leisure–labour trade-off, making them unique to each residence–workplace pair. We decompose the aggregate welfare change in SGE into direct user benefits and wider economic impacts, quantify the statistical uncertainty surrounding each welfare component, and compare the welfare estimates in SGE with those of the mainstream PE method. As a case study, we replicate Greater London with 983 spatial units. We find that the welfare result in SGE is similar in magnitude to the PE CBA outcome in the case of the Elizabeth Line, a major urban rail project. For a series of randomly simulated transport improvements, the two approaches scale proportionately, but the occasional project-specific deviations can be substantial. The paper illustrates that the SGE approach complements aggregate welfare estimates with a detailed spatial pattern of local economic outcomes and is well suited to assessing transport and land-use policies simultaneously.

☆ We thank the editor and three anonymous referees for their in-depth comments and suggestions. We are grateful to Alex Anas, Diego Candia, Tobie Cusson, Matthieu De Lapparent, Davi Guggisberg, Hans Koster, Lukas Makovsky, Dávid Krisztián Nagy, Surabhi Ojha, Chris Parker, Csaba Pogonyi, Daniel Ruiz Palomo, Hugo Silva, Erik Verhoef, and Roger Vickerman for their feedback and technical hints at various stages of this project. Earlier versions and selected parts of this paper have been presented at the 2022 European Conference of the Urban Economics Association (Bocconi University, Milan), the 2023 and 2024 conferences of the International Transport Economics Association (University of Cantabria and University of Leeds), the 2023 Symposium of the European Association for Research in Transportation (ETH, Zürich), the 2023 NLS Colloquium at ETH Zürich, the 2024 Annual Meeting of the US Transportation Research Board (Washington DC), the 2024 European and 2025 North American Meetings of the Society for Benefit-Cost Analysis (University of Warsaw and George Washington University), the 2024 Conference of the Hungarian Society of Economics (Central European University, Budapest), and seminars at Jagiellonian University (Krakow), the UK Department for Transport, the Swiss Federal Railways and the Transport Strategy Centre at Imperial College London. We thank the feedback we received at these events. We take responsibility for any remaining shortcomings. The first author acknowledges funding from the National Research, Development and Innovation Office of Hungary (NKFIH), under Excellence24 grant #151498.

* Corresponding author at: Corvinus Institute for Advanced Studies, Corvinus University of Budapest, Budapest, Hungary.
Email address: d.horcher@uni-corvinus.hu (D. Hörcher).

https://doi.org/10.1016/j.jpubeco.2026.105683
Received 17 February 2025; Received in revised form 9 May 2026; Accepted 25 May 2026

Available online 4 June 2026

1. Introduction

Transport cost-benefit analysis (CBA), often referred to as benefit-cost analysis (BCA) in the US and transport appraisal in Europe, is widely used to evaluate transport investments and guide public spending decisions. In the UK, the EU, and elsewhere, these appraisal frameworks play a central role in assessing large-scale infrastructure projects and in shifting debates on transport investments towards more objective criteria, thereby helping to limit pork-barrel spending (Cadot et al., 2006).

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

The backbone of the mainstream partial equilibrium CBA method is the consumer surplus that a transport improvement brings to existing and new users in the form of travel time savings, mostly computed as a monetary equivalent. We call this quantity the direct transport user benefit (DUB). Since the early 2000s, more research has concentrated on additional welfare gains potentially not captured by travellers' willingness to pay for mobility (Mackie et al., 2011). Wider economic impacts (WEIs) include externalities such as the productivity gain from an increase in access to economic mass. The theoretical foundations and the empirical methods for estimating WEIs are summarised by Venables (2007) and Graham and Gibbons (2019). Note that this classical appraisal framework has remained a partial equilibrium representation of the economy, i.e. it does not capture explicit interactions between the transport market and spillovers in other sectors of the spatial economy.

Despite its widespread use in practice, the partial equilibrium appraisal method has been criticised both within and outside academia. One theoretical concern is a potential overlap between the DUB and WEI layers (see a detailed discussion in Section 2.1). From the viewpoint of practitioners and the wider public, we often witness scepticism because PE cost-benefit analysis seemingly does not capture a wide range of apparent economic implications of a transport improvement, such as job creation and fluctuations in real estate prices. Although microeconomic theory suggests that the partial equilibrium consumer surplus is a reliable measure of economic gain in the absence of market failures in other sectors, there is a widespread preference for evaluation metrics derived from models that transparently replicate observed economic outcomes, such as wages, housing prices, and residential and workplace relocation.

Since the mid-2010s, a new stream of research often referred to as *quantitative spatial economics* (Redding and Rossi-Hansberg, 2017) has made major advances in developing spatial general equilibrium (SGE) frameworks that can be used simultaneously for causal impact evaluation and welfare analysis. Quantitative spatial models (QSMs) are SGE models specifically designed to facilitate parameter inversion and the use of causal techniques in model estimation. They feature two core empirical properties that distinguish them from earlier SGE models.

First, QSMs are invertible, meaning that location-specific parameter vectors (often referred to as locational fundamentals; see Ahlfeldt et al., 2015) can be recovered deterministically from the model under the assumption that the observed data reflect a spatial equilibrium. These structural residuals of the household utility and firm production functions are typically interpreted as residential and workplace amenities and local determinants of firm productivity, although they may also capture measurement error and other unobserved factors. This invertibility property is valuable because quantifying fundamental geographical characteristics in a granular, spatially explicit model is otherwise difficult using conventional empirical methods.

Second, the most important generic parameters of QSMs can be estimated in econometric exercises that remain coherent with the functional forms of the theoretical model. In other words, after transformations that do not violate algebraic equivalence, the equilibrium conditions of the model can be used as econometric estimating equations. The most well-known implication of this property is that the equilibrium condition that describes residence–workplace location choice probabilities turns into a gravity equation after log-linearisation. This enables us to estimate the spread of idiosyncratic preferences in location choice by adopting the econometric toolbox that the trade literature developed for the estimation of gravity equations (Head and Mayer, 2014).

The aim of this paper is to bridge the gap between the state-of-the-art of transport analysis and quantitative spatial economics. A gap exists because mainstream QSMs are based on numerous simplifying assumptions about travel behaviour and transport supply which make their widespread use in contemporary transport analysis challenging. To name the most pressing examples, (i) trip attributes in different dimensions (i.e. monetary cost, travel time, comfort) are not distinguished; (ii) travel disutility is non-linear in travel time and multiplicative, following the *iceberg* (*ad valorem*) specification; (iii) most of the models are based on commuting only, while other trip purposes in passenger transport are ignored; (iv) mode and route choice and congestion are ignored or restricted to functional forms rarely used in traffic modelling; (v) transport supply is exogenous, including decisions on pricing and capacity provision in public transport; and (vi) the welfare predictions of QSMs have yet to be compared with the traditional transport appraisal methodology. We consider these limitations as a research agenda that opens the door to the widespread use of QSMs in transport policy appraisal. In this paper we address items (i), (ii) and (vi) above.

We propose an SGE model in which households face separate temporal and pecuniary budget constraints, with both the time and money costs of travel entering additively into the respective constraints. Unlike in the previous QSM literature, individual labour supply is endogenous, which could provide a basis for future labour economic analyses within the QSM framework. Using the underlying leisure–labour trade-off, we derive and quantify a location-specific expression for the marginal value of time. This specification is not unique in the transport literature; after a logarithmic transformation, our household utility specification becomes very similar to that of Anas and Liu (2007), for example. However, we embed this specification in a framework that is invertible, and the model also allows us to estimate commuting gravity, agglomeration elasticity, and distance decay parameters in a theoretically and geographically coherent manner.

The context of the application of our model is Greater London; the model has nearly one thousand spatial units in line with the Middle-Layer Super Output Areas (MSOAs). We illustrate the model's application in transport policy appraisal through a case study of the Elizabeth Line, a recently completed urban railway line.

The travel time valuations and commuting volumes implied by the model enable us to perform a simplified appraisal of the Elizabeth Line according to the mainstream partial equilibrium CBA methodology. The results are then benchmarked against the estimated net benefit of the project in the SGE model. By adapting recent theoretical advances in Donald et al. (2025) to our framework, we decompose the aggregate welfare change into additive components that can be interpreted as the DUB and WEI quantities in partial equilibrium. In this case, concerns about double counting are eliminated because both components are derived within a single SGE framework. The quantitative SGE approach also enables an important step forward in uncertainty quantification. Since parameter estimation from raw data is part of the computational sequence, we can derive standard errors for each welfare outcome via bootstrapping. In line with prior expectations, we find that WEIs are subject to greater statistical uncertainty than DUBs, and we quantify the magnitude of this uncertainty gap.

Even though the SGE model seems to capture a much richer set of adjustment mechanisms in response to the transport intervention, its aggregate welfare result remains of the same order of magnitude as that of the partial equilibrium CBA. We derive four versions of the partial equilibrium model that differ in (i) whether the value of time varies across origin–destination pairs and (ii) whether new commuting demand induced by residential and workplace relocation is included in the calculation. Under the assumptions of this paper, the SGE welfare result is around 10% lower than the highest partial equilibrium estimate and 31% higher than the lowest. However, the spatial model provides much more detailed insights into the impact of a transport project on the spatial economy, including changes in urban land use, firm productivity, wages, housing prices, and labour supply. The SGE model is also well suited to analysing alternative land-use regulation scenarios associated with the transport improvement. These insights may improve trust in the traditional appraisal methodology and help policymakers and the wider public to better understand the societal impacts of future infrastructure projects.

The remainder of the paper is structured as follows. Section 2 reviews the literature, highlighting the novelties and links to existing work. Section 3 presents the theoretical model. Section 4 describes its quantification. Section 5 outlines the Elizabeth Line counterfactual and its evaluation in partial and general equilibrium. Section 6 concludes.

2. Links with the surrounding literature

This paper relates to two distinct branches of the literature. First, we provide a brief summary of the foundations of cost-benefit analysis in transport. Second, we review previous methods of transport modelling in spatial general equilibrium as well as the emergence of quantitative spatial models, and discuss the features that distinguish the two approaches.

2.1. Transport appraisal in partial equilibrium

Appraisal is one of the most well-documented areas of transport research. Comprehensive reviews were provided by Mackie et al. (2011, 2014); Laird and Venables (2017); Vickerman (2017) and Wangsness et al. (2017).[1]

Economic gains in transport appraisal are usually split into two broad categories. Direct user benefits (DUB) include the economic surplus of transport users, comprising cost savings and the benefits of newly induced travel demand. As a significant fraction of user benefits arises from travel time and discomfort savings, the monetary valuation of time and comfort is a key concern in the development of this literature. The roots of the general theory of time economics date back to Becker (1965) and DeSerpa (1971), while Small and Verhoef (2007); Jara-Díaz (2007), and Small (2012) provide a comprehensive overview of the evolution of thought on time valuation in transport research. Despite the solid theoretical foundations of time use theory, the measurement of the value of time is largely regarded as an empirical matter in practice. Advances in discrete choice modelling enable the estimation of the value of time as a ratio of the marginal utilities of travel time and monetary expenditure using either stated or revealed preference data (for a methodological overview and a large-scale meta-analysis, see Wardman et al., 2016). Recent empirical approaches leverage heterogeneous revealed preferences from large-scale automated datasets, such as metro smart card data (Bansal et al., 2022), toll road data (Bento et al., 2024), ride-hailing data (Buchholz et al., 2025), and a combination of these mode-specific data sources with cellphone traces (Almagro et al., 2024). Even though empirical estimates are sometimes differentiated by income groups or other broad socioeconomic characteristics, the spatial pattern of travel time valuations has not been examined at a high level of granularity.

The second broad category of benefits in transport appraisal is often referred to as wider economic impacts (WEI). WEIs quantify the welfare change associated with market failures outside the transport sector. Relevant market failures include price distortions due to imperfect competition and externalities via agglomeration economies, in other words, changes in economic density (Duranton and Puga, 2004; Mackie et al., 2014; Ahlfeldt and Pietrostefani, 2019). We focus more on agglomeration benefits as this source of WEI is often found to be quantitatively more significant. The theoretical foundations of this approach were laid down by Venables (2007) who showed in a monocentric city model that the firm productivity gains of a transport improvement, stemming from economic density (i.e., city size), are external to the travel cost savings directly perceived by commuters.

Graham and Gibbons (2019) have a detailed review on the estimation of the relevant agglomeration elasticities and their use in the partial equilibrium appraisal model. The standard approach involves four steps: (i) a suitable measure of access to economic mass is defined in function of the transport cost matrix within the city,[2] (ii) in an empirical exercise, the analyst estimates the causal impact on firm productivity (i.e. total factor productivity) of access to economic mass in the form of an elasticity, (iii) using a traffic model, we calculate the change in ATEM induced by the transport improvement under investigation, and (iv) using the estimated agglomeration elasticity, we compute the additional output triggered by the increase in access to economic mass.

Even though in the monocentric city model of Venables (2007), the additional output is clearly external to commuting cost savings, a potential overlap between the DUB and WEI components cannot be unambiguously ruled out. For instance, if firms are themselves transport users, relying on transport services as a production factor, then at least part of the increased output results from a reduction in the price of this factor. Also, when the place of production is modelled as a collection of numerous potential workplaces instead of a unique central business district, then the reallocation of workers between workplaces[3] can fuel user benefits as well as partly overlapping personal gains that Venables' monocentric city model does not capture (Eliasson and Fosgerau, 2019). These unresolved methodological challenges may question the accuracy of surplus calculation in the WEI extension of the partial equilibrium CBA framework. It is widely recognised that a general equilibrium model fully integrating the transport and labour markets would eliminate the risk of double counting. Let us therefore proceed with a review of the related general equilibrium literature.

2.2. Spatial general equilibrium

The wider set of models in which transport is combined with a responsive urban geography is often called land-use transport interaction (LUTI) modelling (see a review in Acheampong and Silva, 2015). This

[1] In most developed and an increasing number of developing countries, the methodology for CBA is published in official guidance manuals, ensuring consistency in the evaluation of projects across different modes of transport and even other sectors. A typical example frequently cited in the academic literature is the UK's Transport Analysis Guidance which is easily accessible through its online platform: see https://www.gov.uk/guidance/transport-analysis-guidance-tag.

[2] Access to economic mass measures the *effective* density of economic activity at a given location. This metric integrates the physical density of the built environment and the scale of firm activity (e.g., employment) with the quality of transport provision. Even if the built environment is characterised by low physical density, fast and affordable connections to other locations via transport infrastructure and services can sustain a high level of effective density.

[3] Reallocation of workers is labelled as the *matching* effect among the three main sources of agglomeration economies in Duranton and Puga (2004).

literature has achieved significant advances in modelling residential, workplace and other activity location choice decisions, but its ability to replicate market mechanisms such as wage and floorspace price formation is limited. Nevertheless, some of the LUTI models, e.g. the one by Hunt and Simmonds (1993), have achieved considerable commercial success and policy impact through their widespread use in ex-ante transport policy evaluation. The narrower literature of spatial computable general equilibrium (SCGE) models features a more rigorous representation of the interlink between transport, labour and housing markets since the pioneering work by Lowry (1964). Extensive reviews are available in Wegener et al. (2011); Robson et al. (2018), and a companion review article of this project by Hörcher and Graham (2025) who survey this literature in light of the latest QSM developments.

A series of studies by Alex Anas and co-authors deserves special attention in this paper. Their urban models developed since the 1990s are usually classified as SCGE. However, their models show many similarities with today's quantitative urban models in the sense that they were pioneers in the adaptation of random utility discrete location choice specifications in spatial equilibrium modelling.[4] On the basis of the discrete choice framework of Anas and Kim (1996), Anas and Liu (2007) published the first spatially explicit SGE model and combined the labour, production and construction markets in urban equilibrium with stochastic traffic assignment. Later on, Anas and co-authors developed further implementations of this model for various case study areas including Los Angeles and Paris (Anas, 2020; Anas and Chang, 2023). A common feature of the Anas-type models is the presence of separate time and money constraints in the household problem, in line with the traditions of the transport literature; see several related references in Footnote 5. Despite many similarities in theoretical structure between this line of research and QSMs, Anas et al. never leveraged the powerful properties of their model for quantitative work (i.e. causal parameter estimation). The model presented in this paper is the first attempt to turn the Anas-type SGE approach into a QSM.

Quantitative spatial models have shifted the prevailing paradigm in spatial equilibrium modelling since the mid-2010s. Section 1 has already provided a broad definition of what makes an SGE model "quantitative": model inversion to recover the vectors of location-specific fundamentals had not been performed previously, and QSMs are also suitable to causally identify some of the core structural parameters in reduced-form econometric models that remain coherent with the specification of the theoretical model. The pioneering urban prototype by Ahlfeldt et al. (2015) was followed by a series of influential contributions by Monte et al. (2018); Heblich et al. (2020); Barwick et al. (2024), and Tsivanidis (2026). QSMs often differ in their geographical scope and assumptions on market structure and the baseline utility and production function. However, this literature places significant emphasis on identifying isomorphisms between the models that differ in the underlying theoretical assumptions. Appendix D of Heblich et al. (2020) provides a valuable methodological overview of the conditions under which five typical SGE specifications, ranging from the urban model of Ahlfeldt et al. to the spatially explicit realisation of New Economic Geography models become isomorphic.

Although transport is a popular subject among the applications of QSMs, this novel literature is still largely disintegrated from the state-of-the-art of transport economic modelling. From one perspective, the authors' experience suggests that QSMs have remained unnoticed by the mainstream transport economics community until recently. From the other perspective, QSMs have several limitations that transport research surpassed a long time ago. We observe new efforts to bridge this gap between the two traditions. For example, Allen and Arkolakis (2022) have incorporated road congestion and stochastic route choice into their previous framework, while Tsivanidis (2026) and Fajgelbaum et al. (2023) integrated mode choice to distinguish car use from other modes of transport. Fajgelbaum and Schaal (2020) applied a QSM to optimise the structure of transport infrastructure networks, while Severen (2023) and Herzog (2024) performed ex-post evaluations of mobility policies. As mentioned in the introduction, several limitations hinder the adaptation of QSMs to transport policy appraisal, one of which is the assumption of *iceberg* transport costs that we relax in this work. Proost and Thisse (2019) pointed to a potential root cause of this research: "The literature of spatial economics has paid too little attention to what has been accomplished in transportation economics (and vice versa)."

3. The model

We present the theoretical model in three steps: after general model definitions, Section 3.1 describes the household problem, Section 3.2 derives choice probability equations, and Section 3.3 introduces the production and floorspace construction sectors of the economy. Appendix A provides further technical details of the model, guiding the reader through each step of the analytical exercise.

Let us consider three groups of agents: households, a production sector and a floorspace construction sector. Geography is represented by a set of discrete locations connected by the transport network. Floorspace in each location is used for both residential and commercial purposes. A model that includes n locations has $n \times n$ possible residence–workplace combinations indexed by ij, where intrazonal commuting with $i = j$ is enabled.

3.1. Household preferences

Urban residents are free to choose where to live and work, and they make location-specific decisions on a continuous scale regarding consumption, residential floorspace use, leisure time, and the intensity of individual labour supply, subject to temporal and monetary constraints. We assume a representative commuter whose location choice probability λ_{ij} is interpreted, based on the law of large numbers, as the share of workers in a population of N individuals who live in i and commute to j.

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

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

In this specification, L_{ij} is a measure of leisure time, K_{ij} is the composite subutility derived from consumption C_{ij} and residential floorspace use H^R_{ij}, γ and β are structural parameters, and z_{ij} is an idiosyncratic taste shock associated with the combination of locations i and j. To keep our notation simple, we suppress the unique identifier of households. Note, however, that z_{ij} takes a different value for each worker. Alternatively, the nested structure in (1) simplifies to the trivariate Cobb-Douglas specification in Eq. (A.1).

Commuters are confronted by two constraints. First, the monetary budget constraint is

$$x_{ij}(w_j - \tau_{ij}) = p_i C_{ij} + q_i H^R_{ij} \qquad [\kappa_{ij}]. \tag{2}$$

In this expression, the wage (w_j), local prices (p_i for consumer goods and q_i for residential floorspace), and commuting cost (τ_{ij}) are exogenous, whereas x_{ij}, the measure of individual labour supply, together with the consumption variables C_{ij} and H^R_{ij}, are endogenous. Second, households also face a time constraint:

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

Here, $\bar{L}$ denotes the daily time endowment. On the right hand side, leisure time L_{ij} and individual labour supply x_{ij} are again endogenous

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

household decisions, while the length of the workday and commuting time t_{ij} are outside the worker's control. In this household problem, κ_{ij} and μ_{ij} are Lagrange multipliers capturing the spatially differentiated marginal utilities of money and time, respectively.

To fully interpret this specification, we emphasise that we consider the workday of the *representative* commuter. Accordingly, we measure individual labour supply x_{ij} on a continuous scale: x_{ij} represents an aggregate measure of the number of workdays supplied by the population, while the daily number of working hours is held fixed. In other words, the model captures the extensive margin of labour supply decisions (participation) rather than intensive responses (hours of work). While restrictive, this choice aligns with empirical evidence (Kleven and Kreiner, 2006; Gutiérrez-i-Puigarnau and van Ommeren, 2010) and with previous transport economic models.[5]

The following derivations are detailed in greater depth in Section A.1 of the Appendix. The first-order condition of the optimal choice of individual labour supply implies

$$\kappa_{ij} w_j = \kappa_{ij}\tau_{ij} + \mu_{ij}(T + t_{ij}), \tag{4}$$

which equates the utility associated with the monetary benefit of the marginal trip to work (on the left-hand side), with the disutility of monetary and time resource requirements (on the right-hand side). Rearrangement leads to an expression of the ratio of the marginal utilities of time and money:

$$\frac{\mu_{ij}}{\kappa_{ij}} = \frac{w_j - \tau_{ij}}{T + t_{ij}} = v_{ij}. \tag{5}$$

We interpret this ratio as a monetary valuation of the incremental relaxation of the worker's time endowment, that is, the (marginal) *value of time* as a resource (DeSerpa, 1971; Jara-Díaz, 2007), denoted by v_{ij}. As one would expect, the worker's wage is among the determinants of the marginal value of time, as foregone time could always be used to earn income via time reallocation to work. Eq. (5) also reveals that the value of time depends on the monetary and time cost of commuting as well. The core consequence from a spatial economic point of view is that the value of time will likely differ between residence–workplace combinations, which is often neglected in mainstream transport policy appraisal. To emphasise this feature of the model, we keep the subscripts of v_{ij} throughout the forthcoming analysis.

Note that v_{ij} can also be interpreted as an effective wage, where the monetary term in the numerator is the daily wage net of the pecuniary cost of commuting, and the denominator normalises this net daily wage by the gross working hours, including the time spent commuting.

Plugging this relationship into the time budget constraint, individual labour supply becomes

$$x_{ij} = \frac{\gamma \bar{L}}{T + t_{ij}}. \tag{6}$$

Individual labour supply increases with the relative importance of non-leisure activities in the worker's time allocation (γ) and decreases with the time requirement of commuting (t_{ij}). This highlights and quantifies a non-trivial channel through which transport supply affects labour markets in the spatial economy, namely the extent to which workers are willing to supply more labour when commuting frictions are reduced.

Our solutions yield the following indirect (sub-)utility functions for a given combination of residential and working locations.

$$\begin{aligned} K_{ij} &= \frac{\gamma \bar{L}\, v_{ij}}{p_i^{\beta} q_i^{1-\beta}} \\ V_{ij} &= \bar{L}\left(\frac{v_{ij}}{p_i^{\beta} q_i^{1-\beta}}\right)^{\gamma} z_{ij} \end{aligned} \tag{7}$$

The effective hourly wage (or in a different interpretation, the marginal value of time), the local price of consumption and residential rents, structural parameters γ and β, and idiosyncratic taste are the only determinants of a residence–workplace combination's attractiveness to households.

3.2. Spatial choice probabilities

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

$$F(z_{ij}) = \exp(-X_i E_j\, z_{ij}^{-\epsilon}), \tag{8}$$

where the average amenity (i.e. the scale parameter) is defined as the product of residence and workplace dependent local fundamentals X_i and E_j, and ϵ governs the spread of individual preferences. These assumptions lead to location choice probabilities that take the form of a commuting gravity equation.[6]

$$\lambda_{ij} = \frac{X_i E_j \left[\frac{v_{ij}}{p_i^{\beta} q_i^{1-\beta}}\right]^{\gamma\epsilon}}{\sum_r \sum_s X_r E_s \left[\frac{v_{rs}}{p_r^{\beta} q_r^{1-\beta}}\right]^{\gamma\epsilon}} \tag{9}$$

The term in the numerator depends on the fundamental attractiveness of i and j and the deterministic part of the indirect utility function in (7). The denominator captures multilateral resistance, i.e. the attractiveness of the alternatives of living and working in i and j. The summation of λ_{ij} over potential workplaces gives the probability of living in i (irrespective of the workplace). Consequently, λ_{ij} divided by the probability of living in i gives workplace choice probabilities conditional on residing in i.

$$\lambda_{ij|i} = \frac{E_j \left(v_{ij}\right)^{\gamma\epsilon}}{\sum_s E_s \left(v_{is}\right)^{\gamma\epsilon}} \tag{10}$$

A key consequence of endogenous labour supply is that employment and population measures are not equivalent. Let the total population be denoted by N. With this, the residential and workplace populations (N_i^R and N_j^W) and the aggregate labour supply by residence and workplace locations (M_i^R and M_j^W) are

$$\begin{aligned} N_i^R &= N \sum_j \lambda_{ij}, & N_j^W &= N \sum_i \lambda_{ij}, \\ M_i^R &= N \sum_j \lambda_{ij} x_{ij}, & M_j^W &= N \sum_i \lambda_{ij} x_{ij}. \end{aligned} \tag{11}$$

The mean individual labour supply in workplace location j is therefore

$$\bar{x}_j = \frac{1}{N_j^W} \sum_i N_i^R x_{ij} = \frac{M_j^W}{N_j^W}. \tag{12}$$

We will return to these features of the model as part of the process of model quantification.

[5] Parry and Bento (2001) and a series of subsequent transport policy appraisal studies built their demand models on the assumptions of (i) endogenous labour supply, (ii) perfect substitution between commuting demand and labour supply, and (iii) the presence of a time constraint in combination with the household's monetary budget constraint. The resulting general equilibrium model provides a transparent expression for the endogenous value of travel time and a tractable link between policy interventions in labour and transport markets. Notable articles following Parry and Bento include De Borger and Van Dender (2003); Arnott (2007); De Borger and Wuyts (2011) and Hörcher et al. (2020). These models remain stylised in a spatial sense. Tikoudis et al. (2015) were the first to adapt this framework to the classical Alonso-Muth-Mills monocentric city model.

[6] See the derivation of the choice probability expression for a multiplicative indirect utility function and a Fréchet-distributed preference shock in the early contributions of the QSM literature, e.g. Appendix Section S.2.3 of Ahlfeldt et al. (2015), or Appendix Section C2 of Heblich et al. (2020).

3.3. Spatial general equilibrium

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

$$Y_j = f^y(A_j, M_j^W, H_j^W) = A_j\left(M_j^W\right)^{\alpha}\left(H_j^W\right)^{1-\alpha} \tag{13}$$

We assume perfect competition in the goods market combined with zero trade cost within the urban area, which allows us to normalise the unit price of the consumption good to $p_i = 1\,\forall i$. Thus, solving the firm's cost minimisation problem and setting marginal cost equal to unity yields the following factor demand functions. (See full derivations in Appendix A.2.)

$$M_j^W = \left(\frac{\alpha A_j}{w_j}\right)^{\frac{1}{1-\alpha}}; \quad H_j^W = \left[\frac{(1-\alpha)A_j}{Q_j}\right]^{1-\alpha} \tag{14}$$

Furthermore, free entry implies zero profits, so that total revenue minus total expenditure on wages and floorspace equals

$$Y_j - w_j M_j^W - Q_j H_j^W = 0. \tag{15}$$

After substituting the production function and (14), straightforward algebraic rearrangement implies that the profit maximising floorspace price is

$$Q_j = (1-\alpha)A_j^{1/(1-\alpha)}\left(\frac{\alpha}{w_j}\right)^{\frac{\alpha}{1-\alpha}}, \tag{16}$$

and the corresponding wage equation is

$$w_j = \alpha A_j^{1/\alpha}\left(\frac{Q_j}{1-\alpha}\right)^{\frac{\alpha-1}{\alpha}}. \tag{17}$$

The general equilibrium setup considers a third group of agents too: the construction sector. Let us define the floorspace production function $H_i = f^h(Z_i, \ell_i)$ as a function of capital Z_i and local land ℓ_i:

$$H_i = f^h(Z_i, \ell_i) = Z_i^{1-\psi}\left(\phi_i(H_i)\cdot \ell_i\right)^{\psi}, \tag{18}$$

where

$$\phi_i(H_i) = 1 - (H_i/\bar{H}_i). \tag{19}$$

With this specification we follow Delventhal and Parkhomenko (2024) who introduce $\phi_i(H_i)$ to capture that local floorspace supply is often constrained by zoning regulations and geographical characteristics. In the specification above, $\bar{H}_i$ represents the theoretical floorspace capacity of each location. When H_i approaches this value, the land requirement for production becomes prohibitively high, meaning that floorspace supply in location i cannot increase any further.

In each location, total floorspace supply is split into residential and commercial uses, $H_i = H_i^R + H_i^W$. The corresponding unit prices may not be identical, $q_i \neq Q_i$, due to differences in tax policies, for instance. We allow such policies to differ between locations, but we assume that the ratio $\xi_i = Q_i/q_i$ remains constant over time.[7] Construction firms take the average floorspace price

$$\bar{q}_i = q_i\left(\frac{H_i^R}{H_i} + \xi_i\frac{H_i^W}{H_i}\right) \tag{20}$$

as given. Similar to the production sector, we assume profit maximisation and perfect competition in the construction sector. After solving the firm's problem, equilibrium floorspace supply becomes

$$H_i = \frac{[(1-\psi)\bar{q}_i]^{(1-\psi)/\psi}\cdot \ell_i}{1 + [(1-\psi)\bar{q}_i]^{(1-\psi)/\psi}\cdot \ell_i/\bar{H}}. \tag{21}$$

This terminates our description of the general equilibrium model's three components: household behaviour, goods production and floorspace construction. Given a set of exogenous parameters, the market clearing conditions above determine the values of endogenous vectors $\{\lambda_{ij}\}$, $\{N_i^R\}$, $\{N_j^W\}$, $\{M_i^R\}$, $\{M_j^W\}$, $\{q_i\}$, $\{Q_j\}$, $\{w_j\}$ and $\{H_i\}$ in spatial general equilibrium.

3.4. Spatial spillovers

Following the empirical literature on agglomeration (Combes and Gobillon, 2015), we decompose local firm productivity into an exogenous component and a multiplier dependent on a measure of economic density, i.e. access to economic mass, denoted by ρ_j.

$$A_j = a_j\,\rho_j^{\eta} \tag{22}$$

Parameter η is the agglomeration elasticity of firm productivity. In terms of the specification of accessibility, the literature offers two alternatives. When the spatial units of the model are large enough to assume that only the density of economic activity *within* zone j affects productivity in j, it is reasonable to consider

$$\rho_j = \frac{M_j}{\mathbb{L}_j}, \tag{23}$$

where M_j is a measure of economic activity, say employment, and $\mathbb{L}_j$ is the geographical land area which may vary between j's. By contrast, when the spatial units are relatively small, it is more likely that productivity spillovers occur between nearby locations. The common specification is then

$$\rho_j = \sum_s \exp(\delta t_{sj})\,M_s, \tag{24}$$

where parameter $\delta < 0$ measures the rate at which agglomeration economies decay over travel time t_{ij}. In this paper we test both alternatives. In addition, we select the effective labour supply as the measure of economic activity, to account for the fact that individual labour supply is endogenous in our model, and therefore employment (i.e., workplace population) alone is not a comprehensive measure of economic mass. In Section 4.3 we use this specification to identify η and the distance decay parameter δ based on the inverted $\{A_j\}$ vector and the observed travel cost matrix and labour supply distribution.

3.5. Welfare analysis

3.5.1. Appraisal in partial equilibrium

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

$$B = N\sum_{ij}\frac{1}{2}\left(\lambda_{ij}^0 x_{ij}^0 + \lambda_{ij}^1 x_{ij}^1\right)\cdot v_{ij}(t_{ij}^1 - t_{ij}^0). \tag{25}$$

This is an approximation of consumer surplus in partial equilibrium, assuming a linear inverse demand curve for travel demand that expresses

[7] The implicit assumption is that policies or external shocks in transport may affect both the residential and commercial floorspace markets simultaneously, but local real estate regulations and thus the price ratio between the two markets remain unchanged. Naturally, this assumption is not suitable when the policy shock itself is a local property regulation that affects residential and commercial floorspace use differently. Such interventions are outside the scope of this paper.

the generalised travel price as a function of travel demand.[8] In Eq. (25), $0.5\,N(\lambda_{ij}^0 x_{ij}^0 + \lambda_{ij}^1 x_{ij}^1)$ is the mean of the commuting volumes in periods 0 (before the intervention) and 1 (after the intervention), and $v_{ij}(t_{ij}^1 - t_{ij}^0)$ is the monetary value of the travel time gain/loss between the two periods. That is, this measure includes both the time savings of those who travelled before the policy intervention as well as the surplus associated with newly induced demand. Note that our model enables us to differentiate the value of time, v_{ij}, between the commuting origin–destination pairs. However, as such differentiation is rarely performed in the state-of-the-practice,[9] we will compute partial equilibrium results with both homogeneous and heterogeneous time valuations.

In the common appraisal exercise we would like to consider the agglomeration externality stemming from improved access to economic mass and the associated productivity gain, in line with the transport appraisal literature (Graham and Gibbons, 2019). The common methodology in this area assumes, based on the urban economic theory developed by Venables (2007), that the additional output that firms realise through improved access to economic mass, keeping all input factors constant, is additional to the direct user benefit of travellers. Using the production function in Eq. (13) and the decomposition of total factor productivity (A_j) in Eq. (22), the output of location j with access to economic mass ρ_j^t is

$$Y_j(\rho_j^t) = a_j\left(\rho_j^t\right)^{\eta} M_j^{\alpha} H_j^{1-\alpha}, \tag{26}$$

where $t \in \{0,1\}$ indicates the equilibria before and after the intervention. The increase in firm output solely attributed to the growth in economic density from ρ_j^0 to ρ_j^1 is

$$\Delta Y_j = Y_j(\rho_j^1) - Y_j(\rho_j^0) = \left[\left(\frac{\rho_j^1}{\rho_j^0}\right)^{\eta} - 1\right] \cdot Y_j(\rho_j^0). \tag{27}$$

We compute this quantity for each location j based on the changes in economic density in response to the Elizabeth Line, as discussed in Section 5.2.1. We then sum the resulting ΔY_j values to obtain an aggregate measure of the agglomeration externality.

Another methodological dilemma we need to address at this stage is whether we consider relocation in the calculation of partial equilibrium (PE) welfare results. Relocation affects both direct user benefits through $\lambda_{ij}^1 x_{ij}^1$ in (25) as well as the external productivity gain through ρ_j^1 in (27). As there is no clear theoretical guidance on whether relocation should be considered in PE welfare analysis, we compute both in the following investigation.[10]

3.5.2. Welfare analysis in spatial general equilibrium

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

Our urban model deviates from the original specification of Donald et al. (2025) in multiple ways. They define a fairly general regional trade model in which consumers live and work in the same location, and utility is additive, with a Gumbel-distributed idiosyncratic shock. Extending location choice to separate residence–workplace pairs is straightforward (see a discussion on this extension in Section 4.6 of Donald et al., 2025). Moving from an additive to a multiplicative specification for household utility may seem more complicated. Proposition 5 of Donald et al. (2025) shows that as long as both the systematic and random parts of the additive utility function are the log transformations of their counterparts in the multiplicative specification, the welfare decomposition below remains transferable.

Our model also differs in the structure of the rest of the spatial economy. While they have a two-tier production sector with intermediate and final goods, our urban model has only one representative producer per location. At the same time, our model features an explicit construction sector which was absent in Donald et al.

Donald et al. (2025) define a pseudo-planning problem whose solution is equivalent to utility-maximising location choices and the corresponding general equilibrium allocation in the spatial model. The Lagrangian of this pseudo-planning problem, after being adapted to our model and notation, is

$$\begin{aligned}
\mathcal{L} = \mathcal{W}&\left(\sum_{ij} N\lambda_{ij}(\Theta)\cdot u_{ij}(L_{ij}, C_{ij}, H_{ij}^R) - \Psi\left(\{\lambda_{ij}(\Theta)\}\right)\right) \\
&+ \sum_{ij} N\lambda_{ij}(\Theta)\cdot x_{ij}\tau_{ij} + \sum_j p_j^{\ell}\cdot \ell_j + \tilde{p}\left[\sum_j A_j f_j^y(M_j^W, H_j^W)\right. \\
&\left. - \sum_{ij} N\lambda_{ij}(\Theta)\cdot C_{ij}\right] + \sum_j \tilde{w}_j\left[\sum_i N\lambda_{ij}(\Theta)\cdot x_{ij} - M_j^W\right] \\
&+ \sum_j \tilde{q}_j\left[f^h(Z_j,\ell_j) - \left(\sum_k N\lambda_{jk}(\Theta)\cdot H_{jk}^R\right) - H_j^W\right] + \sum_j \tilde{p}_j^{\ell}\left[\mathbb{L}_j - \ell_j\right].
\end{aligned} \tag{28}$$

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

$$\Lambda = \mathbb{E}_{ij}\left[p\left(\frac{\partial u_{ij}}{\partial C_{ij}}\right)^{-1}\right]. \tag{29}$$

That is, we normalise household utility by the expected marginal utility of consumption across locations to express its contribution to money-metric social welfare. $\mathcal{W}$ depends on two additive components. The first one is the total systematic utility that households realise through local leisure time, consumption and housing quantities. The second one, $\Psi\left(\{\lambda_{ij}(\Theta)\}\right)$, measures the shift in total utility caused by ex-ante idiosyncratic location preferences under a given distribution of location choices captured by the choice probability vector $\{\lambda_{ij}\}$.[11]

[8] For the theoretical foundations of consumer surplus derived from generalised travel prices, please refer to Glaister (1974).

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

[10] More precisely, this decision in the PE model depends on whether the inverse demand function is defined based on short-run or long-run demand elasticities. In practice, predicting the pattern of household and employment relocation is a non-trivial challenge in the absence of a spatial equilibrium model. In this exercise we assume the partial equilibrium analysis is informed by the SGE model on relocation patterns. The elasticities, ad-hoc gravity models and LUTI models that are frequently used in transport appraisal may lead to different results. In this exercise we avoid this empirical source of bias/incoherence between the PE and SGE methods by passing the SGE outputs over to the PE model.

[11] Section 2.3 and Appendix D of Donald et al. (2025) provide more detail on the functional form of $\Psi\left(\{\lambda_{ij}(\Theta)\}\right)$ under various assumptions on the distribution of idiosyncratic location preferences.

The remaining parts of Eq. (28) are constraints associated with four market clearing conditions. Specifically, $\tilde{p}$, $\{\tilde{w}_j\}$, $\{\tilde{q}_j\}$ and $\{\tilde{p}_j^{\ell}\}$ are Lagrange multipliers corresponding to market clearing in the production, labour, construction and land markets, respectively. Donald et al. (2025) show that in equilibrium, these Lagrange multipliers equal the relevant market prices, so that $\tilde{p} = p$, $\tilde{w}_j = w_j - \tau_{ij}$, $\tilde{q}_j = q_j$ and $\tilde{p}_j^{\ell} = p_j^{\ell}$ for all j. The first constraint states that the amount of goods produced in all j's equals the amount of C_{ij} consumed across residence–workplace pairs. The second constraint equates effective labour supply (including both the extensive and intensive margins) with labour demand. The third constraint equates floorspace production with the sum of residential and commercial floorspace demand in each location. The fourth constraint equates local land supply $\mathbb{L}_j$ and demand ℓ_j.

Consider a set of transport policy interventions that induce small changes in the travel time matrix $\{t_{ij}\}$. Based on Eq. (28) and the envelope theorem, the associated change in aggregate welfare is

$$\begin{aligned} dW = & \underbrace{\sum_{ij} \Lambda\mu_{ij} \cdot N_{ij}x_{ij} \cdot (-dt_{ij})}_{\text{(i) time savings}} \\ & + \underbrace{N \cdot \mathrm{Cov}(\Lambda - \kappa_{ij}^{-1}, \kappa_{ij} p \cdot dC_{ij})}_{\text{(ii) dispersion of MU of consumption}} + \underbrace{N \cdot \mathrm{Cov}(\Lambda - \kappa_{ij}^{-1}, \kappa_{ij} q_i \cdot dH_{ij}^R)}_{\text{(iii) dispersion of MU of housing}} \\ & + \underbrace{\sum_{ij} \tau_{ij} \cdot d(N_{ij}x_{ij})}_{\text{(iv) fare revenues}} + \underbrace{\sum_j \ell_j \cdot dp_j^{\ell}}_{\text{(v) land value}} + \underbrace{\sum_j \frac{p\eta Y_j}{\rho_j} d\rho_j}_{\text{(vi) agglom. externality}} . \end{aligned} \tag{30}$$

Item (i) transforms time savings into social welfare, using the origin–destination-specific marginal utility of time (μ_{ij}) and the average inverse marginal utility of money that we defined in Eq. (29). Items (ii) and (iii) originate from the unequal spatial distribution of the marginal utilities (MU) of consumption and housing. The sign and magnitude of these terms depend on the covariance of two terms: an ij-pair's marginal utility of money relative to the city-wide average and its change in consumption and housing demand in response to the policy. For example, if dC_{ij} (or dH_{ij}^R) is positive in those residence-workplace pairs where the marginal utility of money is also higher than average, then items (ii) and (iii) are expected to make a positive contribution to the aggregate welfare change. Component (iv) captures fare revenues; even though this cash flow is a transfer within society, it must be taken into account with a positive sign because commuter utility is expressed net of the fare payment. Item (v) is land value uplift; this effect is part of the welfare change because land revenues are retained by absentee landlords as an isolated part of social surplus. Finally, item (vi) is the externality that firms generate when they become better connected and thus more productive.

In Section 5 we use the sum in (30) as the measure of the aggregate welfare change and quantify its six constituents as well, for a detailed comparison with partial equilibrium appraisal. The intermediate steps that lead from (28) to (30) are reported in Appendix D.

Other papers in the QSE literature rely on an alternative approach to compute the change in aggregate welfare. Similar to the logsum formula in logit discrete choice models, there is a closed-form expression for the expected value of household utility in the discrete choice framework that features a Fréchet-distributed random utility shock.[12]

$$\mathbb{E}[U_{ij}] = \Gamma\left(\frac{\varepsilon - 1}{\varepsilon}\right) \left[\sum_{i,j} X_i E_j \left(\frac{v_{ij}}{p_i^{\beta} q_i^{1-\beta}}\right)^{\gamma\epsilon}\right]^{1/\epsilon} \qquad \forall i,j \tag{31}$$

In this expression $\Gamma(\cdot)$ is the Gamma function. Note that in spatial equilibrium, expected utility is the same across all residence–workplace combinations but the marginal utility of money is not. To derive a monetary measure of the economic gain of households in this model, we compute the compensating wage that achieves the same improvement in expected utility as the transport project. The compensating wage is then multiplied by N, the number of workers. As the production and construction markets are perfectly competitive, economic surplus is not generated there. The only element that we need to add to the monetised change in expected utility is the predicted change in land revenues, $\sum_j \ell_j \cdot dp_j^{\ell}$, which also appeared in the welfare decomposition of (30).

4. Data and model estimation

The structure of the theoretical model enables the quantification of two types of exogenous parameters. First, this model features five vectors of location-specific parameters that capture geographical or local regulatory characteristics: residential and workplace amenities X_i and E_j, local productivity A_j, and two vectors capturing local land and real estate regulation, ξ_i and $\bar{H}_i$. Second, the model has a set of spatially undifferentiated structural parameters: ϵ, α, β, γ, ψ, $\bar{L}$, T, and two parameters relating to agglomeration economies to be introduced soon. We discuss the estimation/calibration of these two sets of parameters in this section.

The numerical implementation of the model covers 983 spatial units defined by the UK Office for National Statistics as middle layer super output areas (MSOA) in Greater London. The paper relies on publicly available data sources only; with access to the Secure Research Service of the ONS, the model could also be implemented at a higher level of spatial granularity, including the LSOA or even the OA statistical units.[13] With the current spatial resolution, the model distinguishes $983^2 = 966{,}289$ potential residence–workplace combinations in location choice. Our data sources are listed below.

(i) Commuting matrix: ONS 2011 Census. The decennial Census of the UK includes large-scale datasets on the residential and workplace locations of the population. The publicly available version of this dataset is aggregated to the MSOA level. Given that the 2021 Census has been unconventional in many ways due to the Covid-19 pandemic, the remaining data sources in this paper were collected to complement the 2011 Census.

(ii) Residential and workplace populations: ONS 2011 Census. In principle, the row and column sums of the commuting matrix correspond to the residential and workplace populations of employed individuals for each spatial unit. The total residential and workplace populations, including both active and inactive individuals, are also available on the ONS website. While we will not use this data in the model implementation, comparing the total population distribution with the commuting matrix offers additional descriptive insights for the analysis.

(iii) Wages: Annual Survey of Hours and Earnings, Table E.2. We use the net daily income column of this publicly available dataset, provided at an MSOA level for each year, including 2011, as an approximation of local wages.

(iv) Floorspace prices at an LSOA level (aggregated into MSOAs): Our baseline data source for floorspace prices is the open-source dataset of Ahlfeldt et al. (2023). They match the real estate sales data of the Land Registry with Energy Performance Certificate data, and estimate the floorspace price index using a mix of parametric and non-parametric estimation techniques. Note that this dataset does not include a breakdown of floorspace prices into commercial and residential properties. Therefore, we use two

[12] For the detailed derivation of this expression, see Appendix Section S.2.2 of Ahlfeldt et al. (2015) or Appendix Section C1 of Heblich et al. (2020).

[13] See a detailed description of statistical geographies on the ONS website: https://www.ons.gov.uk/methodology/geography/ukgeographies/statisticalgeographies.

further data sources described as items (v) and (vi) to infer the wedge between commercial and residential floorspace prices (ξ_i).

(v) Commercial floorspace prices at Local Authority District level: UK Valuation Office Agency administrative data. This dataset as well as item (vi) are only available at a borough level in London; we apply inverse distance weighting to interpolate this data to the MSOA level.

(vi) Residential floorspace prices at Local Authority District level: Tenant Services Authority RSR.

(vii) Transport infrastructure networks: OpenStreetMap (OSM). We use the entire road and pedestrian path network of Greater London to infer walk times to the nearest public transport stop.

(viii) Bus network and timetables: Bus Open Data, General Transit Feed Specification (GTFS): Transport for London.

(ix) Railway and metro network and timetables, GTFS: Rail Delivery Group. This dataset contains information on the service provider of each rail and metro line, which enables us to remove certain lines (i.e., the Elizabeth Line) in counterfactual scenarios using the *gtfstools* package in R.

(x) Historic population densities from the online data appendix of Heblich et al. (2020), with the corresponding author's permission.

The numerical implementation of the model is preceded by a series of data processing steps. The travel times and monetary travel costs between MSOA centroids, that is, the t_{ij} and τ_{ij} matrices, are computed via the April 2024 version of the routing package called *r5r*; see Pereira et al. (2021). This library combines infrastructure shapefiles from OSM with public transport timetables stored in GTFS format. The routing algorithm considers walking distance to the nearest stop/station, in-vehicle travel times according to the timetable, and transfer times, to produce a realistic estimate of the travel time and monetary cost matrix along the shortest path.[14] This approach also allowed us to compute both matrices with and without the Elizabeth Line, providing the main input to the counterfactual simulations in Section 5.2.

When quantifying the model, we exploit the recursive structure of the quantitative spatial model and proceed through the following steps: first, we estimate the gravity equation implied by the location choice model, which yields an estimate of ϵ, the Fréchet shape parameter. Second, we use the estimated $\hat{\epsilon}$ to invert the model and recover three sets of location fundamentals: amenities X_i and E_j, and the productivity vector A_j. Model inversion is a deterministic process: we apply the equilibrium conditions of the model to express one location-specific vector as a function of observed and previously identified parameters. In other words, the model structure enables us to establish a one-to-one relationship between the observed data and the unknown vectors of location fundamentals, and quantify the latter as structural residuals of the model. Third, based on the resulting vector of local productivities, we estimate the elasticity of productivity with respect to agglomeration (η) and the associated distance decay (δ). Fourth, we perform two more deterministic steps to recover the exogenous part of the local productivities and the floorspace construction limits $\bar{H}_i$. Standard errors are calculated using a bootstrap procedure with 250 replications of the estimation sequence.

The remaining structural parameters are not separately estimated at the current phase of this research project. We borrow the following values from the literature: $\alpha = 0.8$ from Valentinyi and Herrendorf (2008); $\beta = 0.75$ from Davis and Ortalo-Magné (2011); $\psi = 0.25$ from Combes et al. (2019); we set $\bar{L} = 24$ hours and $T = 8$ hours by intuition.

[14] More specifically, we compute the monetary cost of travel by metro on the basis of the origin and destination station, while the cost of bus trips depends on how many boardings are necessary along the shortest path. We apply 2021 tariff levels. Under these assumptions, the cost of a daily round-trip is 1.82 percent of the daily wage on average (median: 1.66 percent; max: 5.10 percent).

4.1. Commuting gravity

Model specification. Model quantification begins by estimating ϵ, the spread of the Fréchet-distributed idiosyncratic shock in households' utility function. Following Ahlfeldt et al. (2015), we express the location choice probability equation in (9) as

$$\log N_{ij} = \alpha_0 + \vartheta_i + \vartheta_j + \nu \cdot \log \upsilon_{ij} + \varepsilon_{ij}, \tag{32}$$

In the equation above, N_{ij} is the number of observed travellers in the commuting matrix, α_0 is the intercept capturing the denominator (multilateral resistance) in gravity Eq. (9), ϑ_i is a residence (origin) fixed effect, ϑ_j is a workplace (destination) fixed effect, $\nu = \gamma\epsilon$ becomes the coefficient of bilateral resistance $\log \upsilon_{ij}$, and ε_{ij} is the error term. Note that υ_{ij} in Eq. (5) is expressed as a function of wages (w_j), the monetary and time costs of commuting (τ_{ij} and t_{ij}), and the number of working hours per day (T), which we treat as a constant. Thus, υ_{ij} in the estimating equation is derived deterministically from the data. The main parameter of interest, ν, is transferable between this estimating equation and the location choice probability function in Eq. (9).

The functional form of (32) is common in the broader empirical literature on gravity estimation in international trade (Head and Mayer, 2014). This allows us to apply a robust and widely adopted estimator from that literature, the Poisson Pseudo-Maximum Likelihood (PPML) method of Santos Silva and Tenreyro (2006). The use of this model is motivated by two concerns. First, N_{ij} on the left-hand side of (32) includes many zeros in the commuting matrix, as the commuting flow is effectively zero between two-thirds of the MSOA pairs in our data. In an OLS log-log model, these observations would have to be removed, which would imply a substantial loss of information. The second concern arises from Jensen's inequality: under heteroskedasticity, the parameters of a log-linearised model estimated by OLS lead to biased estimates because $E[\log y] \neq \log E[y]$.

Identification. Identification of ν in (32) is challenging because bilateral resistance may be endogenous: unobserved factors can affect both resistance and flows. For example, travel times may be affected by the non-random placement of infrastructure, and road congestion may give rise to reverse causality. Thus, the estimated elasticity in (32) should not automatically be given a causal interpretation. Accordingly, our empirical strategy relies on fixed effects and the use of Euclidean distance between i and j as an instrumental variable to isolate plausibly exogenous changes in bilateral resistance. Under standard identifying assumptions, the conditional regression function in (32) can then be interpreted as approximating the average response of flows to changes in resistance.

Unfortunately, due to the relatively recent opening of the Elizabeth Line in 2022, which also coincided with the aftermath of the Covid-19 pandemic and followed the latest Decennial Census in 2011, we are unable to combine the single cross-section of pre-intervention commuting data with a commuting matrix reflecting the post-intervention equilibrium. Panel commuting flow data would allow us to weaken our exclusion restriction substantially. Given this data limitation, we follow the majority of the literature by assuming that changes in travel times perfectly determine changes in commuting flows. By contrast, the panel gravity approach applied by Severen (2023); Herzog (2024) and Lee and Tan (2024) would enable us to include pair-specific fixed effects to control for the time-invariant, hard-to-measure determinants of commuting flows, while letting origin and destination fixed effects be time-varying, capturing the non-commuting effects of the policy.

Results. The results for six model specifications are summarised in Table 1. Models (1) and (2) are OLS estimates of the model, with and without origin and destination fixed effects. These models rely on a restricted sample because OLS cannot handle zero flows after the log transformation. It is remarkable though that the elasticity estimate in model (2) is close to our most preferred one in (4). Models (3) and

Table 1
Estimating commuting gravity.

	Dependent variable: log commuting flow					
	(1)	(2)	(3)	(4)	(5)	(6)
Impedance	value of time (effective wage)				travel time	gen.cost
Notation	v_{ij}	v_{ij}	v_{ij}	v_{ij}	t_{ij}	$\bar{v}t_{ij} + \tau_{ij}$
Method Instrument	OLS†	OLS†	PPML	PPML + IV Eucl.dist	PPML	PPML
Estimate	2.959***	12.303***	19.035***	10.193***	−0.893***	−0.855***
	(0.012)	(0.021)	(0.059)	(0.031)	(0.008)	(0.008)
Fixed effects	no	yes	yes	yes	yes	yes
RMSE	1.078	0.657	8.815	8.684	13.948	13.686
AIC	1,052,942	708,040	2,767,492	2,836,697	5,290,953	5,524,191
BIC	1,052,974	729,218	2,790,654	2,859,856	5,314,115	5,547,353
# of obs.	352,300	352,300	966,289	965,306	966,122	966,122

† Only origin-destination pairs with positive flows are included.
Standard errors in parentheses, *** 99%, ** 95%, * 90%.

(4) are Poisson models with fixed effects. A common endogeneity concern is that impedance between i and j may not be independent of the flows themselves. One potential reason may be the non-random placement of infrastructure. To address this concern, we instrument v_{ij} with the Euclidean distance in model (4), which provides our preferred estimate.

Assuming that commuters have a fixed travel time budget of one hour on average (Kung et al., 2014), such that $\gamma = (T+1)/24$, $\epsilon = \nu/\gamma$ is computed from our preferred empirical estimate of $\hat{\nu} = 10.193$. Note that all four elasticities and the resulting $\hat{\epsilon}$ values are higher in models (2) to (4) than the single-digit estimates in other QSM studies in the literature.

In models (5) and (6), we estimate the gravity equation using travel times and generalised travel costs as impedance measures. This exercise shows that our data yield a negative commuting elasticity of similar magnitude to that reported in mainstream QSMs when the standard impedance measure is used. The higher Fréchet shape parameter in our main model can be interpreted as less randomness in people's location choice decisions. We attribute this outcome to the fact that v_{ij} is a more comprehensive measure of the trip generating forces in the commuting context. That is, it has a higher explanatory power in commuting decisions than the pure travel time.

4.2. Amenity residuals

Exploiting the recursive structure of the model, in this step we rely on the estimated $\hat{\epsilon}$ parameter to recover the values of the residential and workplace amenity vectors $\{X_i\}$ and $\{E_j\}$. Based on the definitions and location choice probabilities defined above, the link between workplace and residential populations can be expressed as

$$N_j^W = \sum_i \lambda_{ij|i} \cdot N_i^R. \tag{33}$$

We can express fundamental amenity levels E_j after substituting (10) into (33):

$$E_j = N_j^W \left(\sum_i \frac{v_{ij}^{\gamma\epsilon} N_i^R}{\sum_s E_s v_{is}^{\gamma\epsilon}} \right)^{-1} \tag{34}$$

This implies an equation for each location in function of all other E_j's that we can solve for iteratively to recover the unique vector of workplace amenity levels from the observed distribution of residential and workplace populations and the observed determinants of v_{ij}. Residential amenities are recoverable in a similar fashion. Let us introduce

$$\tilde{X}_i = X_i \left(q_i^{1-\beta} \right)^{-\gamma\epsilon}, \tag{35}$$

which captures all residence-dependent endogenous variables of location choice probability (9). The values of $\tilde{X}_i$ are the solution of

$$\tilde{X}_i = N_i^R \left(\sum_j \frac{v_{ij}^{\gamma\epsilon} N_j^W}{\sum_r \tilde{X}_r v_{rj}^{\gamma\epsilon}} \right)^{-1}. \tag{36}$$

From the solutions, X_i is quantified using data on residential floorspace prices q_i and by inverting (35) above. Given the choice probability equation in (9), the amenity vectors can be scaled freely without affecting spatial outcomes and welfare predictions. We normalise $\{X_i\}$ and $\{E_j\}$ such that their geometric means equal one.

A visual representation of the residuals in Fig. 1 reveals insightful patterns. The results show that the locations in Central London, with the exception of Westminster and the nearby MSOAs, are not particularly pleasant places to work, conditioned on the very high wages offered by these places. An inner ring surrounding the City of London and Canary Wharf has particularly low amenities for working. By contrast, some of the locations in the suburbs are more attractive for working than what their wage levels would justify. Lower density and congestion externalities (i.e. noise, pollution) may explain some of these results. By contrast, the most appealing residential locations are clustered around the central areas of Finchley Road and Swiss Cottage, Kensington and Chelsea, and the low-density residential neighbourhoods of Richmond.[15]

In Appendix C.2, we estimate the elasticity of the amenity residuals recovered in this stage with respect to various measures of economic mass. As the magnitude and even the sign of the resulting elasticities are highly sensitive to the choice of mass measure, we decided not to include these experimental results in the core analysis. Accordingly, residential and workplace amenities remain exogenous in the policy application of the model.

4.3. Agglomeration externalities

The vector of local productivities $\{A_j\}$, or total factor productivity in the production function (15), is computed by rearranging the wage equation in (17).

$$A_j = \left(\frac{w_j}{\alpha} \right)^{\alpha} \left(\frac{Q_j}{1-\alpha} \right)^{1-\alpha} \tag{37}$$

As we observe both wages (w_j) and commercial floorspace prices (Q_j), the A_j vector is computed by directly substituting the observed data into (37). Our empirical objective is to isolate the exogenous, geographically determined component of productivity (a_j) from the impact of economic density (ρ_j). Based on the productivity decomposition (22), the empirical objective is to estimate agglomeration elasticity $\eta = \partial \log \hat{A}_j / \partial \log \rho_j(\delta)$, in which $\hat{A}_j$ are the productivity residuals recovered via (37). The core endogeneity concerns well-known in the literature are that (i) density may be correlated with unobserved local characteristics, e.g. natural advantages/endowments, that also affect productivity, and (ii) through the endogenous location choice of firms and competition forces, productivity may also affect the magnitude of agglomeration, fueling reverse causality (Combes and Gobillon, 2015; Graham and Gibbons, 2019). To identify the causal effect of agglomeration on productivity, we deploy instrumental variables and the control function technique in three alternative specifications.

Specification and identification. The models reported in Table 2 differ in the specification of ρ_j, the inclusion of fixed effects, and the method of estimation. In models (1) and (2), the general functional form in (22) remains unchanged but we ignore agglomeration spillovers between locations, so the specification of ρ_j follows (23).

[15] Even though, from a technical perspective, the vector of residuals we quantify via model inversion may include noise and a wide range of hard-to-measure determinants of local characteristics, the pattern of high values coincides with what is typically perceived as the most pleasant neighbourhoods of London, especially in residential terms. Accordingly, we find it reasonable to refer to these residual vectors as amenities.

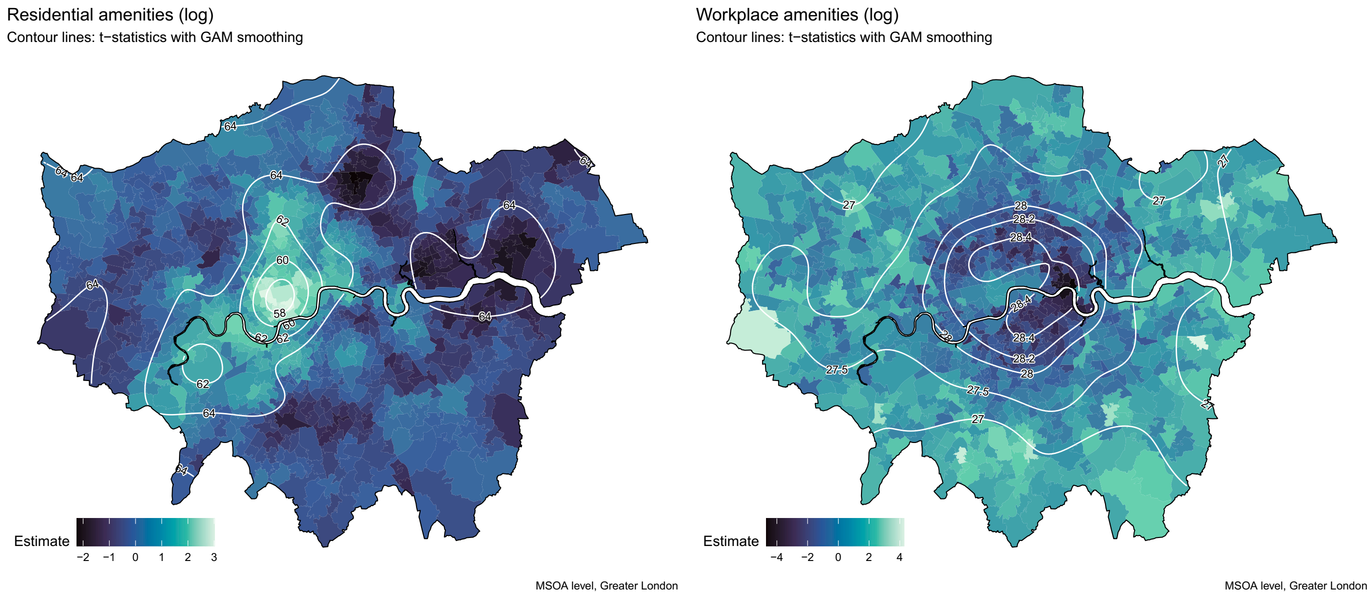

Fig. 1. The spatial distribution of residuals, interpreted as workplace and residential amenities and recovered via model inversion. Contour lines represent smoothed surfaces of t statistics of local amenities computed through bootstrapping.

Table 2
Agglomeration economies and distance decay.

	Dependent variable: log productivity residual			
	(1)	(2)	(3)	(4)
Density measure	Emp.density	Emp.density	Total emp.†	Total emp.†
Method	OLS	2SLS	NLS + CF	NLS + CF
Productivity elasticity, η	0.125***	0.115***	0.150***	0.142***
	(0.005)	(0.009)	(0.015)	(0.024)
Distance decay, δ			−0.080***	−0.071**
			(0.012)	(0.012)
Borough fixed effects	no	yes	yes	yes
RMSE	0.17	0.11	0.06	0.06
AIC	−672.44	−4251.30	−2697.48	−2597.99
BIC	−657.77	−4085.01	−2404.04	−2304.55
# of obs.	983	983	983	983

† Total employment is aggregated in 2.5-min travel time bands.
Standard errors in parentheses, *** 99%, ** 95%, * 90%.

Model (1) is a crude OLS estimate. In model (2) we instrument $\log \rho_j$ with a set of historical and geographical variables. Specifically, we use population densities in 99 Greater London boroughs in 1861 and 1921.[16] Historical population concentrations are predictive of current density, particularly in London, where much of the historical building stock has been preserved. At the same time, conditional on contemporary controls, historical population patterns are plausibly exogenous to present-day productivity shocks in London's finance, innovation, and service oriented economy. We also include a third-order polynomial of the logarithm of distance from the City of London, the central business district. The underlying assumption is that distance from the CBD is correlated with density, particularly in monocentric cities, due to the fundamental trade-off between commuting costs and competition for land. However, this measure of relative geographical location is assumed not to affect firm productivity through mechanisms other than economic density.

Models (3) and (4) have been estimated in a two-step process which captures the decay in spillover effects between nearby MSOAs, following Koster (2024). In the first step we create 2.5-minute-wide concentric travel time doughnuts around location j denoted by $\mathcal{R}_{jr}$, aggregate the effective labour supply in each ring and estimate the contribution of each distance ring to the measure of density in j via the parameter vector $\{d_r\}$. Thus, the estimating equation is specified as

$$\log \hat{A}_j = \eta \log \sum_r \left(d_r \sum_{k \in \mathcal{R}_{jr}} N_k^W \bar{x}_k \right) + \vartheta_{j \in z} + \varepsilon_j. \tag{38}$$

As $\log \hat{A}_j$ depends on η and the ring-specific parameters non-linearly, we estimate the model with nonlinear least squares (NLS). Then, in the second step we fit a curve on the coefficient estimates to quantify δ as the parameter of τ_r, the mean travel time between locations in ring r and location j:

$$\log \hat{d}_r = \delta \tau_r + \varepsilon_r. \tag{39}$$

We follow Koster (2024) again by applying a control function approach. As an initial step, we estimate the nonlinear model in (38) with NLS, ignoring that the resulting $\hat{d}_r$ estimates are likely biased. We use these estimates to compute $\hat{\rho}_j$, the access to economic mass predicted by the model. Then, in a first-stage regression we regress this outcome using exogenous historical ATEM measures and distance from the CBD. In model (3) of Table 2, the control function is

$$\log \hat{\rho}_j = \sum_{y \in \mathcal{Y}} \log \rho_j^y + f(\mathrm{dist}_j) + \varrho_j, \tag{40}$$

where $\mathcal{Y}$ includes the years 1841, 1861, 1881, 1901 and 1921, ρ_j^y is the population density of location j in year y, $f(\mathrm{dist}_j)$ is the third-order polynomial of the distance from the central business district, and ϱ_j is the error term. Model (4) differs from model (3) in the assumed control function: we use 30-minute travel time doughnuts and a nonlinear specification similar to the second-stage regression as an instrument of today's ATEM.

$$\log \hat{\rho}_j = \sum_{y \in \mathcal{Y}} \eta^y \cdot \log \sum_{r=1}^{3} \left(d_r^y \sum_{k \in \mathcal{R}_{jr}^y} \mathrm{pop}_k^y \right) + \varrho_j, \tag{41}$$

[16] Our data source, with the corresponding author's permission, is the online appendix of Heblich et al. (2020) who used original data provided by the Cambridge Group for the History of Population and Social Structure (Wrigley, 2011).

We take the 1861 and 1921 population distributions and travel time matrices from Heblich et al. (2020) to estimate this model with nonlinear least squares.

The residual vectors of (40) and (41), which we denote $\hat{\rho}_j$, are correlated with the dependent variable and, thus, also correlated with ε_j in the main regression because economic density is endogenous. Thus, when we recover the $\hat{\rho}_j$ vector and insert it into the second-stage regression

$$\log \hat{A}_j = \eta \log \sum_r \left(d_r \sum_{k \in R_{jr}} N_k^W \bar{x}_k \right) + \vartheta_{j\in z} + f(\hat{\rho}_j) + \tilde{\varepsilon}_j, \tag{42}$$

the original error term is decomposed into $\varepsilon_j = f(\hat{\rho}_j) + \tilde{\varepsilon}_j$. As $f(\hat{\rho}_j) = E[\varepsilon_j|\rho_j]$ is the conditional mean of ε_j, from which $\tilde{\varepsilon}_j$ is independent by construction, we ensure that the remaining error $\tilde{\varepsilon}_j$ is no longer correlated with the endogenous ATEM measure. In other words, we correct for endogeneity and the resulting η and d_r estimates are unbiased.[17] In practice, we use the third-order polynomial of $\hat{\rho}_j$ as $f(\hat{\rho}_j)$.

Note that the estimation of the agglomeration parameters is part of the sequence of empirical work stages detailed in this section. There is inherent uncertainty in $\hat{A}_j$, the dependent variable of the present exercise. Consequently, the standard errors produced by the OLS, 2SLS, and NLS estimation algorithms are not informative. The standard errors reported in Table 2 are computed by bootstrapping the entire model inversion and parameter estimation process 250 times.

Results. Let us turn to the empirical results in Table 2. The naïve model in column (1) yields an agglomeration elasticity of 12.5%. The 2SLS regression with employment density as the density measure leads to a higher result. The two control function specifications lead to estimates between the OLS and 2SLS results. Our preferred model is (4), with an elasticity of $\hat{\eta} = 14.2\%$ and a distance decay of $\hat{\delta} = -0.071$, which implies that spillovers fade quickly after 15 to 20 min of travel time. Note that some of the estimated $\hat{d}_r$ coefficients in (42) may not be statistically significantly different from zero. Fig. C.1 in the Appendix plots these estimates together with the resulting distance decay curve for model (4); δ has been estimated using only the $\hat{d}_r$'s that we found statistically significant at the 95% confidence level.

The estimated agglomeration elasticity $\hat{\eta} = 14.2\%$ is at the higher end of the values found in the previous literature (see Graham and Gibbons, 2019). However, our result does not stand out from previous studies focusing on Greater London, specifically. Dericks and Koster (2021) developed a QSM using the same case study context. Exploiting the exogenous variation in density caused by bombings during WW2, the agglomeration elasticity they estimate is even higher than ours, 19.6%. To explore potential factors behind the high elasticity we found, we re-estimated the less computationally intensive model (3) using MSOAs located under and above the median distance from the CBD separately; see Table C.1 in the Appendix. We find that the elasticity rises to 17.5% for the inner-city subsample while it declines to 8.1% for the more peripheral parts of London. As we move from the centre to the city edge, this clearly non-linear pattern seems to converge close to the UK-wide average of 4.3% recommended by the UK Transport Appraisal Guidance and the 4.7% mean of the global literature reported by Graham and Gibbons (2019).

4.4. Productivity fundamentals and density limits

Given the above estimated $\hat{\eta}$ and $\hat{\delta}$ parameters, the observed travel time matrix, and the quantified values of A_j, we recover the $\{\hat{a}_j\}$ vector from Eqs. (22) and (24).

$$\hat{a}_j = \hat{A}_j \left[\sum_s \exp(\hat{\delta} t_{sj})\, N_s^W \bar{x}_s \right]^{1/\hat{\eta}} \tag{43}$$

The $\{\hat{a}_j\}$ vector captures fundamental local characteristics that make firms more productive in certain locations, controlling for the level of access to economic mass.

Fig. 2 provides a visual illustration of the decomposition outlined in Eq. (22). The patterns suggest that, when it comes to the geographical determinants of productivity, the East and South of Greater London are particularly disadvantaged, and the triangle between the City of London, Camden and Kensington is the most well-endowed part of the city. The right-hand side of the figure shows that economic density significantly contributes to productivity. Its relative multiplier effect can exceed a factor of two: ceteris paribus, access to economic mass makes well-connected locations more than twice as productive as the least connected ones. It is, of course, not surprising that the places with the highest *exogenous* productivity are also the most centrally located ones. The theory of urban spatial structure suggests that cities typically develop around the most productive locations, creating a path dependency in urban development.

Finally, $\xi_i = Q_i/q_i$ is calculated as the ratio of observed commercial and residential floorspace prices while the local density limit, $\bar{H}_i$, is recovered by inverting (21).

$$\bar{H}_i = \frac{\Xi_i(\bar{q}_i, \xi_i, L_i)}{\Xi_i(\bar{q}_i, \xi_i, L_i)\cdot (H_i^R + H_i^W)^{-1} - 1}, \qquad \text{where} \quad \Xi_i = \left[(1-\psi)\bar{q}_i\right]^{\frac{1-\psi}{\psi}} L_i. \tag{44}$$

We assume that the local fundamentals X_i, E_j, a_j, ξ_i and $\bar{H}_i$ remain constant in counterfactual policy simulations.

5. Policy appraisal

In this section we apply the fully quantified model to perform an illustrative policy appraisal experiment. The infrastructure policy we selected as a case study is the implementation of the Elizabeth Line, a large-scale urban rail investment programme that has attracted a lot of attention in the academic literature as well. The counterfactual application of the model assumes that the Elizabeth Line is introduced in the baseline 2011 spatial equilibrium of Greater London; we let the model converge to a new spatial equilibrium, and assess the economic impact of the new intervention in a comparative statics exercise. Among the main outputs of the simulation, we present predictions on the spatial redistribution of economic activity, including market outcomes such as the wage and floorspace price distribution (see Section 5.2), and an estimate of the aggregate welfare of the policy in Section 5.3.

Note that this paper does not aim to deliver a comprehensive cost-benefit analysis of the Elizabeth Line. As the concluding section of the paper discusses in more detail, our model does not account for several important sources of welfare gains and losses, such as potential reductions in congestion and pollution externalities. That said, our aggregate welfare estimates are not directly comparable to the official appraisal reports on the Elizabeth Line. However, in Section 5.3, we do perform a comparative analysis to learn more about how the spatial equilibrium welfare effect relates to the results of a comparable partial equilibrium CBA. This methodology-oriented experiment represents a first step towards reconciling the QSM approach to transport CBA with the mainstream PE framework.

5.1. The spatial distribution of time valuations

Before turning to policy simulations, let us visualise an outcome of the calibration process that we devote increased attention to. Fig. 3 plots the spatial distribution of v_{ij}, our theoretical measure of the marginal

[17] Recall though that the predicted $\hat{\rho}_j$ values in the initial estimation of (40) and (41) are based on biased $\hat{d}_r$ values. Thus, we create an iterative process and use the $\hat{d}_r$ parameters estimated in the second-stage regression in (42) to recalculate $\hat{\rho}_j$, and repeat the first- and second-stage regressions until convergence is reached. In practice, with a three-digit tolerance, convergence has been reached after just the second iteration.

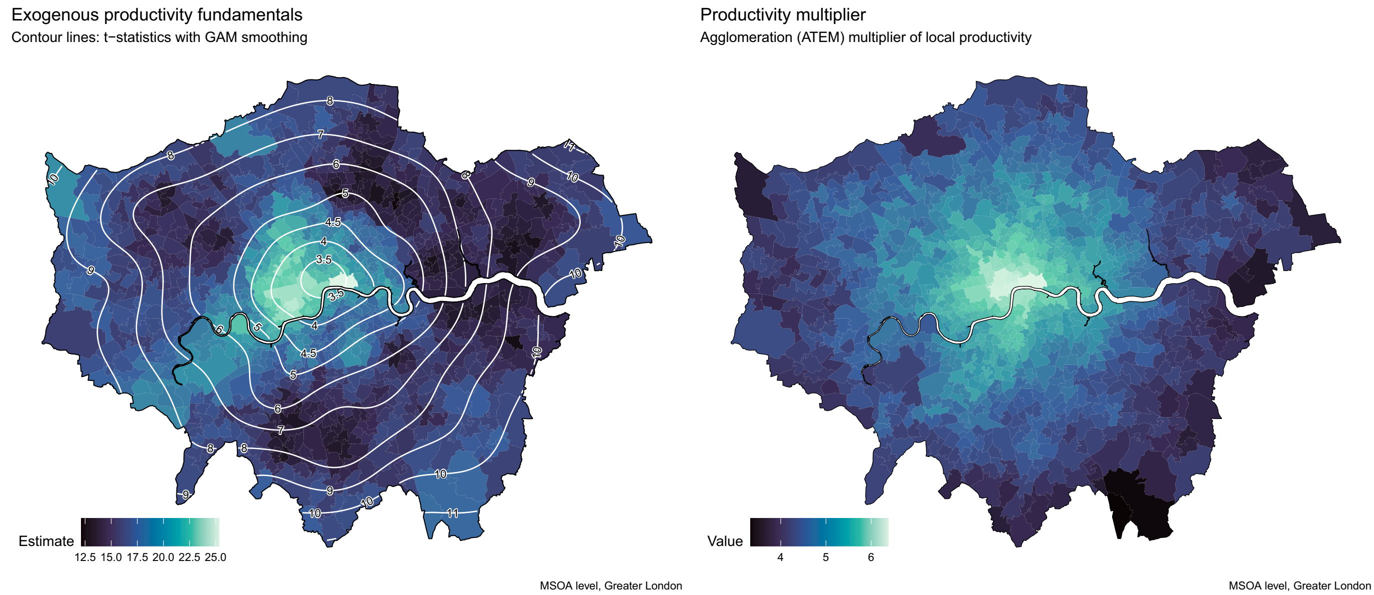

Fig. 2. Productivity as a locational fundamental and the multiplier generated by agglomeration.

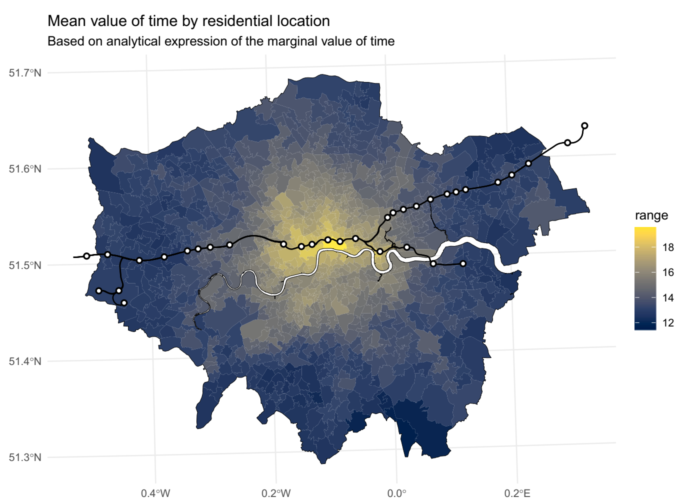

Fig. 3. Spatial distribution of mean travel time valuations, measured in 2011 GBP/hour and averaged by residential location.

value of time. In this plot we depict their mean by residential location. As one would anticipate, residential locations in central London have the highest mean value of time, given the concentration of high-wage individuals whose commuting distance is also likely relatively small. Overall, time valuations vary between around 12 and 20 GBP/hour. The unweighted mean value of time among all OD-pairs is around 11.6 GBP/hour, while the flow- and distance-weighted means are 14.5 and 14.4 GBP/hour, respectively. These averages are broadly consistent with the empirical estimates of the value of time from the early 2010s.[18]

This exercise offers a unique perspective on the spatial distribution of travel time valuations. Such insight is unprecedented, mainly because empirically measuring the value of time using traditional stated or revealed preference data at a comparable level of spatial disaggregation would require an excessively costly data collection effort. By contrast, our value of time distribution is derived from publicly available and easily accessible wage and commuting time/cost data. We believe that the geographical disaggregation of travel time valuations can be useful from a practical perspective, in both demand modelling and policy evaluation.

5.2. A large-scale infrastructure intervention

The fully calibrated QSM is now ready to perform counterfactual simulations; that is, to modify some of the exogenous parameters of the model to mimic a policy intervention and estimate its impact on urban form and general equilibrium outcomes.

In this exercise we assume the hypothetical scenario that the Elizabeth Line opens in the 2011 state of the London economy, leading the city to transition to a new spatial equilibrium. The descriptive plot in Fig. 4 shows the impact of the policy on the distribution of the mean travel time from individual residential locations (left side) and to individual workplaces (right side). To make the spatial pattern of time savings more visually perceptible, we apply quantile-based colour breaks in these plots, as indicated by the colour spectra in the legend. In the most privileged residential location, the representative commuter gains 12.51 min in travel time. Naturally, the gain is the highest along the alignment of the new railway line, but some of the northern MSOAs, from which residents may transfer to the Elizabeth Line, also gain a non-negligible amount. The mean travel time savings that workplaces experience, many of which are located in well-connected places already, vary within a tighter range.

In the counterfactual simulations we assume that the transport improvement is combined with a relaxation of the exogenous constraint on floorspace development in Central London. We increase the value of $\bar{H}$ by 2% within a 1.5 km radius of five stations: Paddington, Bond Street, Tottenham Court Road, Farringdon and Liverpool Street. This assumption is meant to capture the impact of explicit changes in local zoning regulations as well as a tendency to grant discretionary permission for real estate development more readily near Elizabeth Line stations.[19] In

[18] Note again that v_{ij} in our model is the pure opportunity cost of leisure time, which does not include other components of the value of travel time savings relating to reliability, comfort, crowding, etc.

[19] A study by GVA (2018) documents that the Elizabeth Line was associated with formal planning approvals that enabled commercial intensification near key stations. It finds that "over the 2008–16 period 48% of permitted planning applications within a 1 km radius of Elizabeth Line stations have made direct

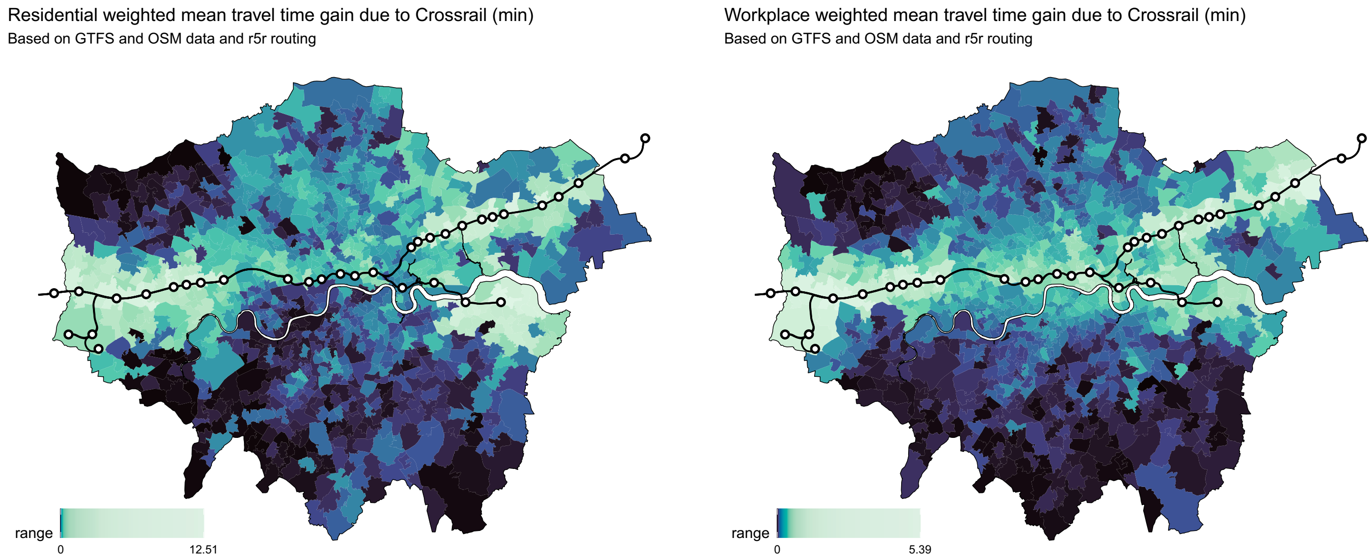

Fig. 4. Policy simulation: mean travel time savings provided by the Elizabeth Line for a representative resident (left) and worker (right).

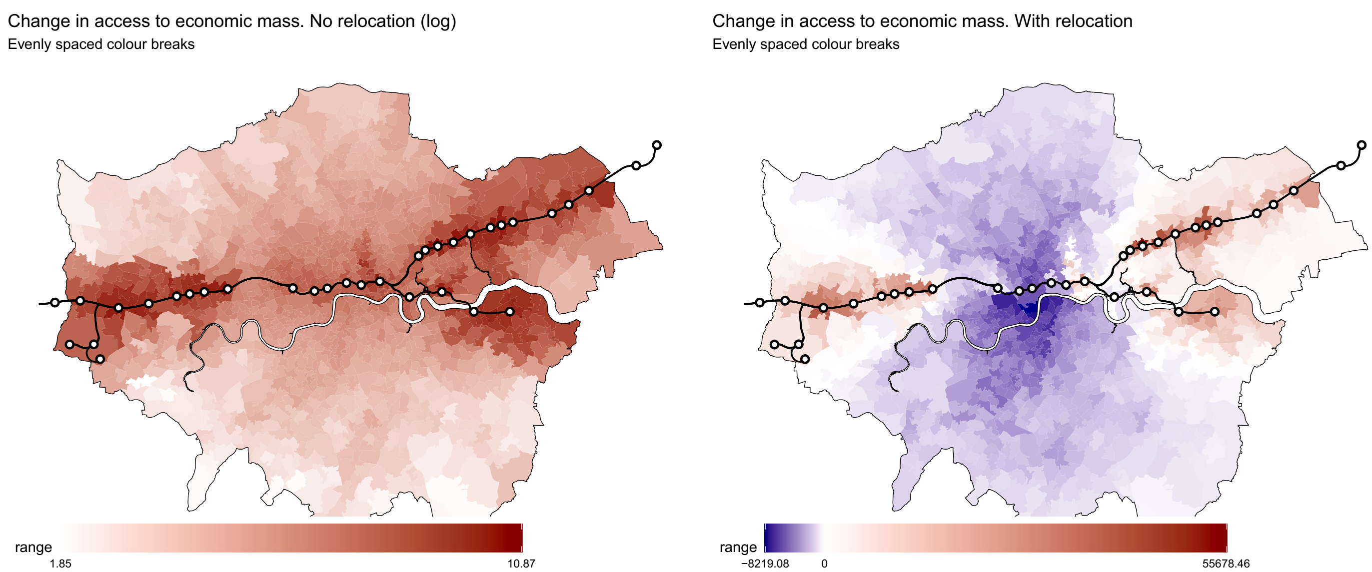

Fig. 5. Counterfactual outcomes: Changes in economic density (access to economic mass) with and without relocation.

Appendix E we test the sensitivity of our results with respect to this ad-hoc assumption.

5.2.1. Economic density, firm productivity and wages

A key way in which transport improvements influence the urban economy is by increasing the effective economic density of the locations they connect. In other words, the policy reduces the impedance between firms (business-to-business connectivity) and firms and households (labour pooling), thereby fostering agglomeration economies. Based on our decomposition of local firm productivity in Eqs. (22) and (24), let us show the transformation of effective density in two steps.

First, the left-hand side of Fig. 5 shows the change in access to economic mass in response to the new travel time matrix t_{sj} only, keeping the distribution of employment N_s^W and individual labour supply $\bar{x}_s$ constant. As the travel time effect of the transport project is strictly negative,[20] this adjustment in economic density is always positive. In line with intuition, this effect is greatest near the new railway stations, especially in the suburban parts of the line which may have been more poorly connected before the intervention.

Second, the right-hand side of Fig. 5 plots the change in access to economic mass after households and firms relocate in spatial equilibrium. In this case we allow all endogenous variables, including prices and quantities in the labour and housing markets, to readjust to the new spatial equilibrium. That is, we allow both the masses and spatial impedances

reference to Crossrail supporting their proposals in their planning application documents," and specifically names Liverpool Street, Bond Street, Farringdon, and Tottenham Court Road as places where office development peaked (page 32).

[20] Note that in this model we ignore congestion. In the presence of congestion spillovers in the network, transport costs may increase on certain origin–destination pairs.

to readjust. The primary insight from the plot is that the change in economic density is indeed not positive everywhere. In fact, a significant part of the city, including the suburbs in the North and South as well as the Centre of London experiences a reduction in economic density. Note from the legend of this figure though that the positive changes in economic density are much greater in magnitude than the negative changes. Also, the positive changes are highly skewed due to a few outliers that experience extremely high gains in access to economic mass.

The fact that employment somewhat decreases in the centre of London is a less anticipated outcome of the simulation. We attribute this result to the tightness of the real estate market in the initial equilibrium; that is, to the extremely high level of commercial floorspace prices in the CBD. As the new rail line makes the less space-constrained locations in the suburbs much better connected, their attractiveness relative to the CBD increases, which may incentivise employers in Central London to relocate along the Elizabeth Line. However, we would like to emphasise that the model does not predict an economic downturn in Central London: later we will see that both the wage rate and the amount of commercial floorspace supply actually increase in and around the City of London.

Fig. 6 illustrates the difference between the two sides of Fig. 5, which reflects the impact of firm relocation (and the readjustment of individual labour supply) on access to economic mass. In other words, it highlights the consequences of shifting from a partial equilibrium model to a spatial general equilibrium framework. The negative and positive deviations are more evenly distributed in this case. We infer from the resulting pattern that Central London's effective economic density might be overestimated in a partial equilibrium model while the agglomeration gain of the suburban workplaces along the Elizabeth Line is underestimated when relocation is ignored in PE. Fig. E.1 in the Appendix plots the change in density without relocation against the change in density with relocation, at the level of individual MSOAs. It shows that many places gain slightly more when the effect of relocation is taken into account. However, in a significant number of cases, we find in Fig. E.1 that the model predicts an increase in economic density in the absence of relocation, while it actually decreases when employment relocation is unlocked.

How do the changes in access to economic mass translate into effects on firm productivity? Fig. 7 depicts the local changes in total factor productivity A_j. Note again that we use a quantile-based colour scheme which allows us to make extreme values and subtle differences between small changes visible at the same time. The right-hand panel of Fig. 7 confirms that the distribution of the changes plotted here has a long tail in the positive direction. The most affected firms are concentrated at a small number of new stations along the Elizabeth Line: the neighbourhoods that gain the most are Hillingdon and Ealing in West London and Redbridge, Newham, Havering, Bexley, Greenwich and Barking and Dagenham in the East. The rest of the city experienced very little change in firm productivity. Regarding the main business districts, the model predicts a negligible reduction in productivity in the City of London ($\Delta A_j = -0.056$) and a negligible improvement in Canary Wharf ($\Delta A_j = +0.08$). Note, however, that our urban model does not consider long-distance connectivity. As one of the main purposes of the the Elizabeth Line is to make the City of London and Canary Wharf better connected with Heathrow Airport in the West and City Airport in the East, this urban model may not provide a comprehensive assessment of firm productivity.

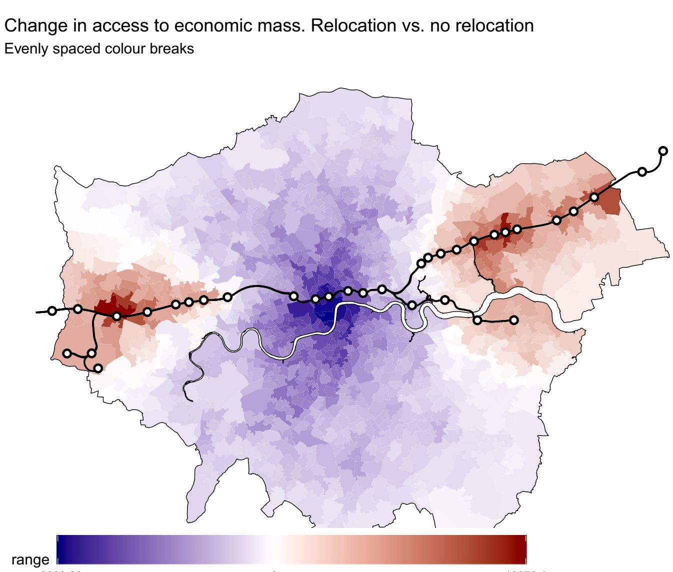

Fig. 6. The impact of employment relocation on access to economic mass: the difference between the predicted changes in economic density with and without relocation in Fig. 5.

5.2.2. Relocation, land-use, floorspace prices and relocation

The consequences of the productivity boost in terms of wages and floorspace prices, plotted in Fig. 8, are clear and intuitive. The wage rate grows along the entire length of the Elizabeth Line by up to GBP 8 per day at the most privileged locations. By contrast, remote areas in the South and the North may even lose in terms of wages due to their loss of relative attractiveness in terms of connectivity; this is often called an agglomeration shadow (Redding and Sturm, 2008). Wages mildly increase in the biggest employment centres: by GBP 0.23 per day in the City, GBP 0.30 in Canary Wharf, and GBP 0.22 in Westminster. When it comes to floorspace prices, the model predicts mostly positive changes. We do not observe a significant increase in floorspace prices in Central London, where prices are already the highest and the floorspace restriction is relaxed by 2% according to the assumed policy scenario.

The residential and workplace relocations are plotted in Fig. 9. By construction of the model, relocation is governed by the new wage and floorspace price distributions as well as commuting times. Since both the wage and floorspace price changes exhibit similar patterns (with a correlation of 77%), we do not have clear prior expectations regarding relocation patterns. From Fig. 9 we infer that the suburban stations of the Elizabeth Line attract new workplaces mostly, at the expense of the residential population. Fig. 10 on the redistribution of floorspace supply confirms that there is tight competition between the residential and commercial uses of land: residents are relocating from the MSOAs where the highest increase in commercial floorspace demand emerges.

Looking more closely at Fig. 9, the destination of the residents priced out from the proximity of the Elizabeth Line is less clear. Minor negative and positive changes in the residential population are spread sporadically across North and South London. Notably, the dark blue bands near the Elizabeth Line, where the residential population has decreased significantly, are surrounded by dark red areas. This suggests that residents cannot compete with new commercial floorspace users in the most attractive locations but are instead relocating as close as possible, forming rings around the newly emerging employment centres along the Elizabeth Line.

In summary, we observe a pattern of employment decentralisation combined with higher wages, floorspace prices and floorspace density in the CBD of London. The general pattern of residential floorspace supply points towards a tendency of suburbanisation, but not to the immediate proximity of the new infrastructure. Heblich et al. (2020) documented in a century-long context that the 19th-century development of London's suburban rail network enabled residents to separate their residential and workplace locations. That rule continues to hold in the 21st century. Although the mean commuting time decreases by 0.3%, the mean Euclidean commuting distance *increases* by 1.96%, rising from 11.72 to 11.95 kilometres.

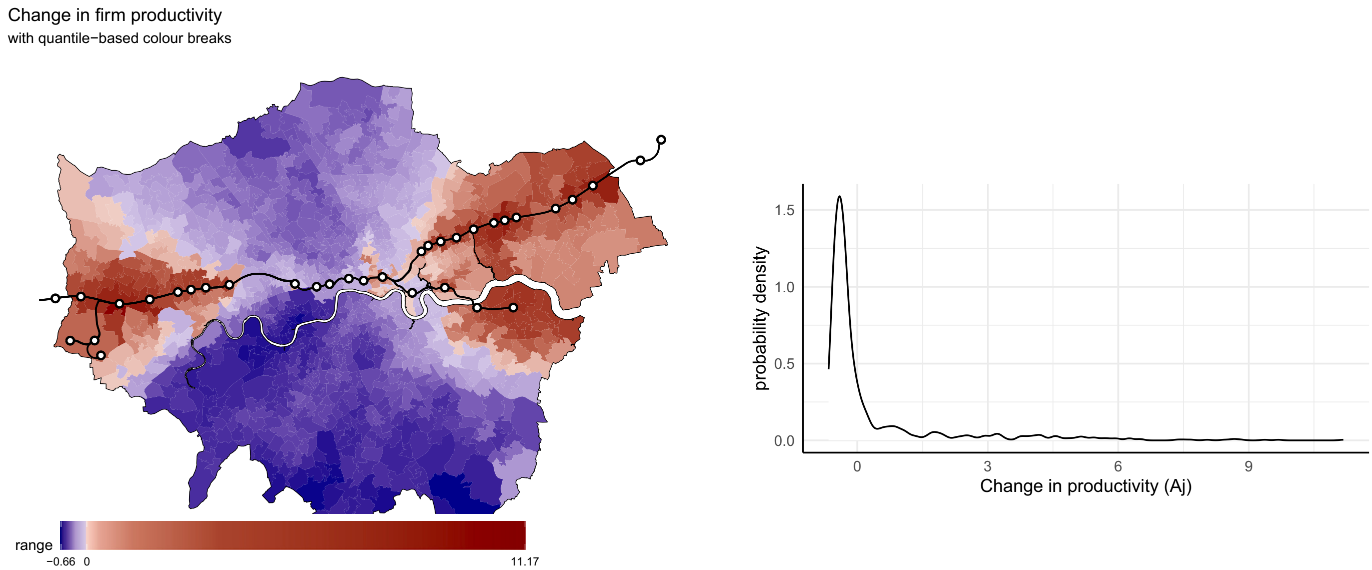

Fig. 7. Counterfactual outcomes: firm productivity.

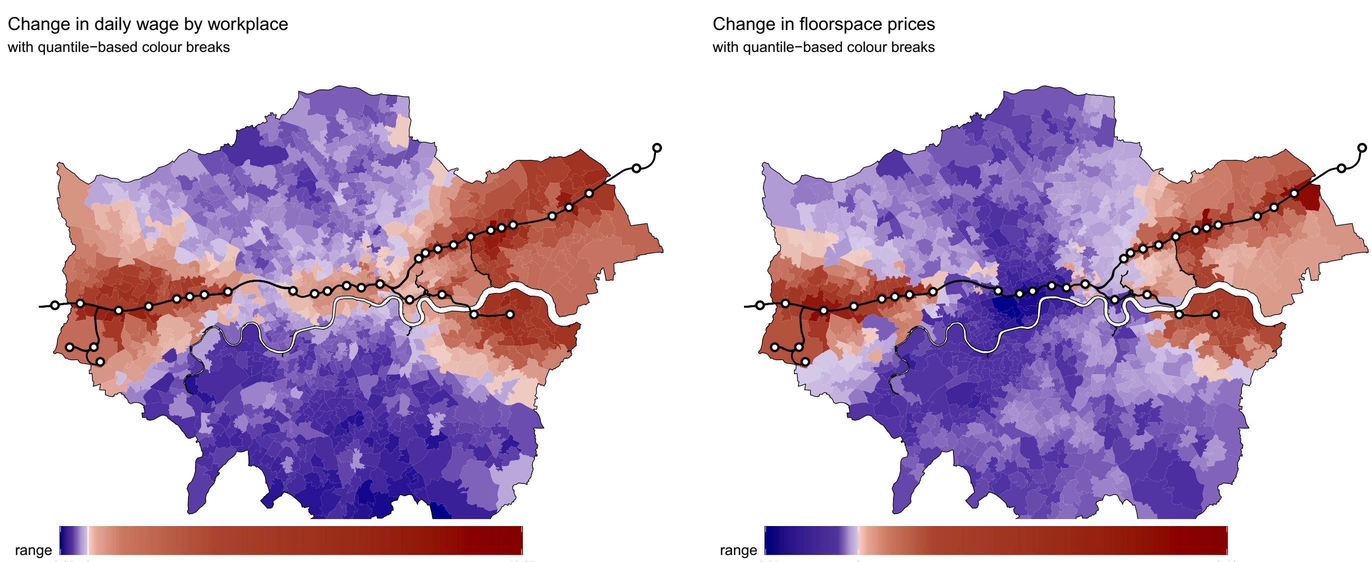

Fig. 8. Counterfactual outcomes: wages and floorspace prices.

5.2.3. *Further spatial insights*

Finally, we would like to draw attention to some of the less conventional spatial outcomes of a transport intervention, plotted in the two panels of Fig. 11. These capture the impact of the policy on (i) travel time valuations, which are unique to this paper in the QSM literature, and (ii) on the mean income by workplace locations, taking the wage rate as well as individual labour supply into consideration.

The left-hand-side of Fig. 11 reveals significant changes in the distribution of travel time valuations, that is, the distribution previously shown in Fig. 3. This plot illustrates the fact that the transport improvement itself affects how commuters value time. We observe a pattern of increased travel time valuations in suburban locations. This might be explained by the relocation of high-wage commuters to these zones, thus increasing v_{ij} in the model. In some cases, the changes predicted by the model are clearly substantial, ranging up to GBP 1.46 per hour. This finding is remarkable in light of the fact that the city-wide mean value of time is nearly unaffected by the policy: it increases by GBP 0.015 or 0.1% only. Apparently, the city-wide average hides significant spatial variation which may have important implications locally. Once again, these are previously unseen insights in the context of transport appraisal that normally keeps the value of time constant, both spatially and temporally.

The right side of Fig. 11 shows the change in mean income by workplace, i.e. the product of the wage rate and the individual labour supply: $w_j \bar{x}_j$. This outcome is interesting to compare with the changes in productivity (in Fig. 7) and the pure wage rate (in Fig. 8), which indicate significant heterogeneity. When individual labour supply is taken into account, Fig. 11 shows that income per worker increases everywhere along the Elizabeth Line. That is, it is unlikely that any commuters employed at workplaces along the new infrastructure would experience a reduction in individual income as a result of the intervention.

5.3. *Alternative appraisal approaches*

In the third step of the analysis we study the use of the SGE model developed in this paper in a welfare economic cost-benefit analysis, that

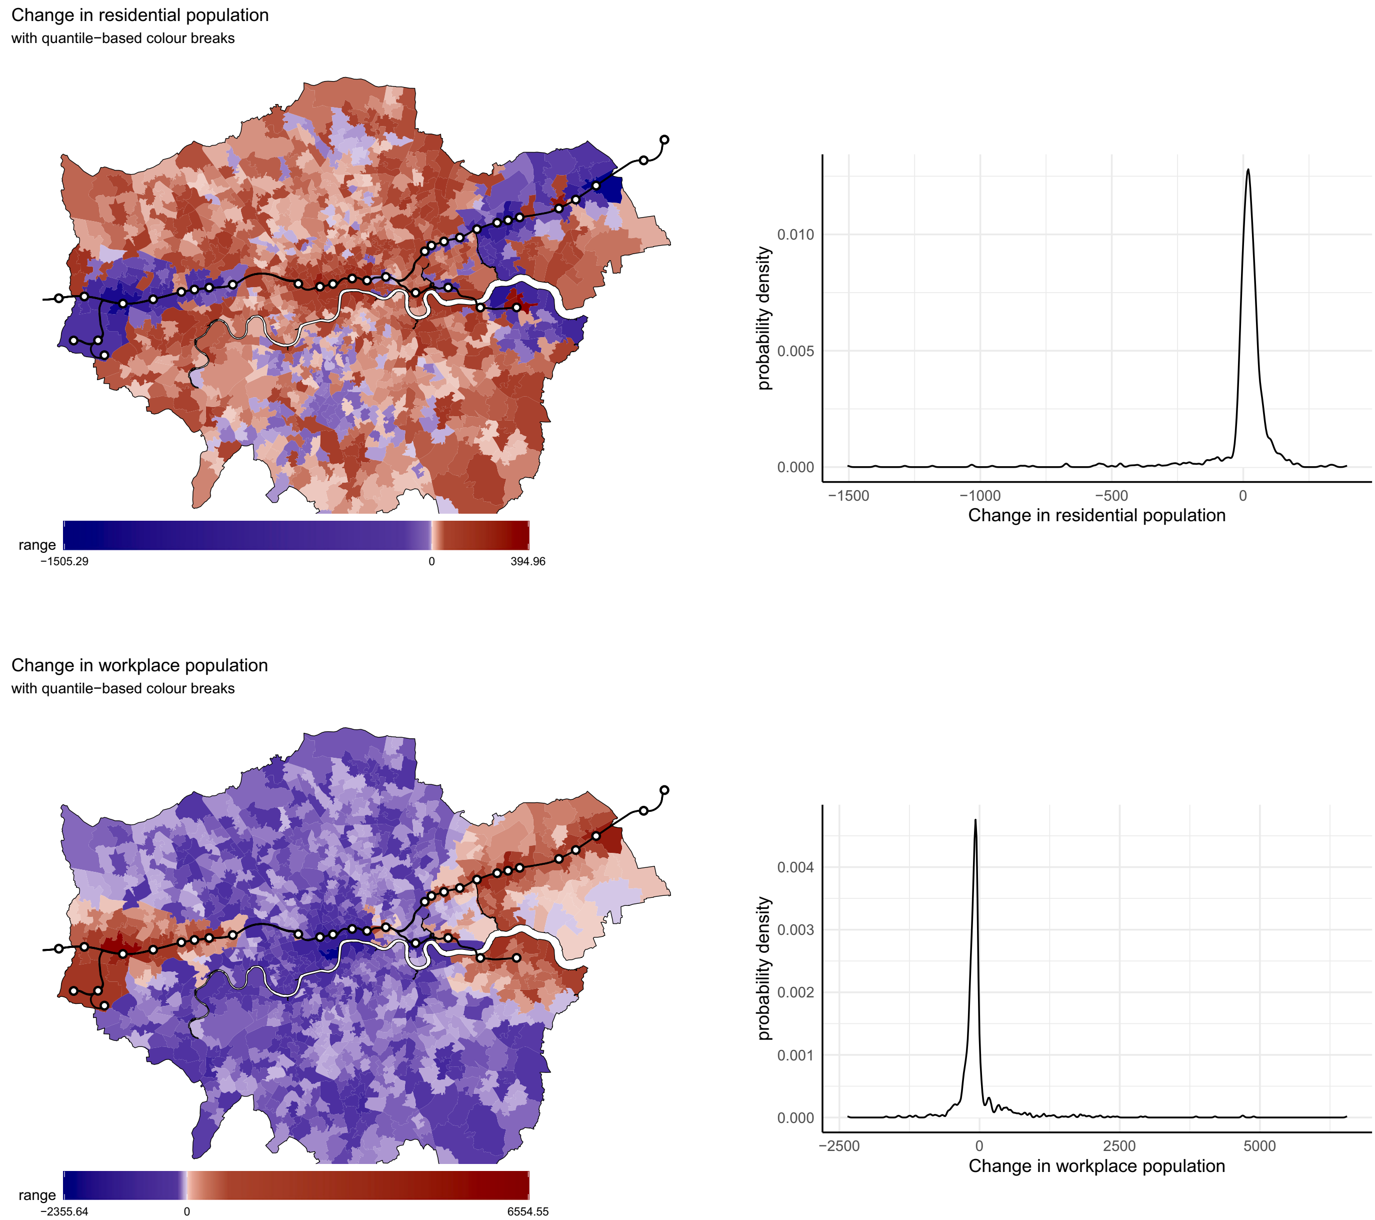

Fig. 9. Counterfactual outcomes: predicted relocation patterns in response to the Elizabeth Line.

is, a transport appraisal exercise. Section 2.1 concluded that the state-of-the-practice in transport appraisal is limited to a partial equilibrium representation of the transport market, and most of the benefits are derived from the monetary valuation of the aggregate travel time savings delivered by an intervention. The core research question of this subsection is whether the QSM approach developed in this paper is expected to produce fundamentally different evaluation results in comparison with the mainstream PE methodology.

The results of our numerical experiment are presented in Table 3. The table reports both the aggregate welfare change (final row) and an additive disaggregation of its components. We also make use of the fact that parameter estimation is an integral part of the appraisal model. This allows us to quantify uncertainty at the level of each welfare component by bootstrapping the entire sequence of estimation and simulation stages and deriving standard errors and significance levels numerically. Uncertainty quantification substantially improves our understanding of the confidence we can assign to various sources of economic gains.

Some of the results shed light on less known properties of partial equilibrium transport appraisal as well. First, we find that the direct user benefit is considerably higher when the mean of the value of time is applied instead of locally differentiated ones. We find 10 and 15 percent higher direct user benefits in the former case in the PE models without and with relocation, respectively. Second, we observe a complementary relationship between the direct benefit and wider economic impact layers: as we enable relocation in the third and fourth columns, the volume of direct user benefits increases due to induced travel demand, but the value of wider economic impacts shrinks (see Fig. 6 and the related discussion in Section 5.2.1). These complementary effects partly neutralise each other. It is more concerning that the estimated productivity gain in the PE model *with* relocation is subject to substantial statistical uncertainty. With the standard error recovered via bootstrapping, the estimate does not pass even the 90% confidence level. This implies that, in addition to the theoretical inconsistencies of the partial equilibrium welfare analysis with relocation, we have serious concerns about the statistical robustness of the productivity gain computed in this way.

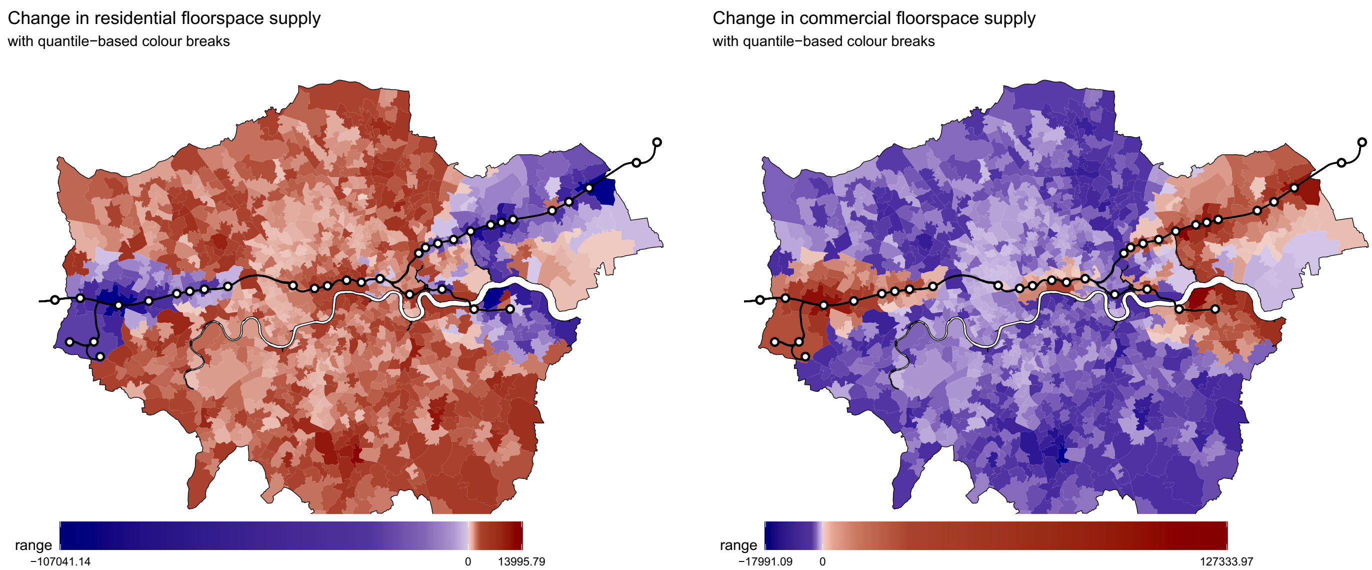

Fig. 10. Counterfactual outcomes: floorspace supply.

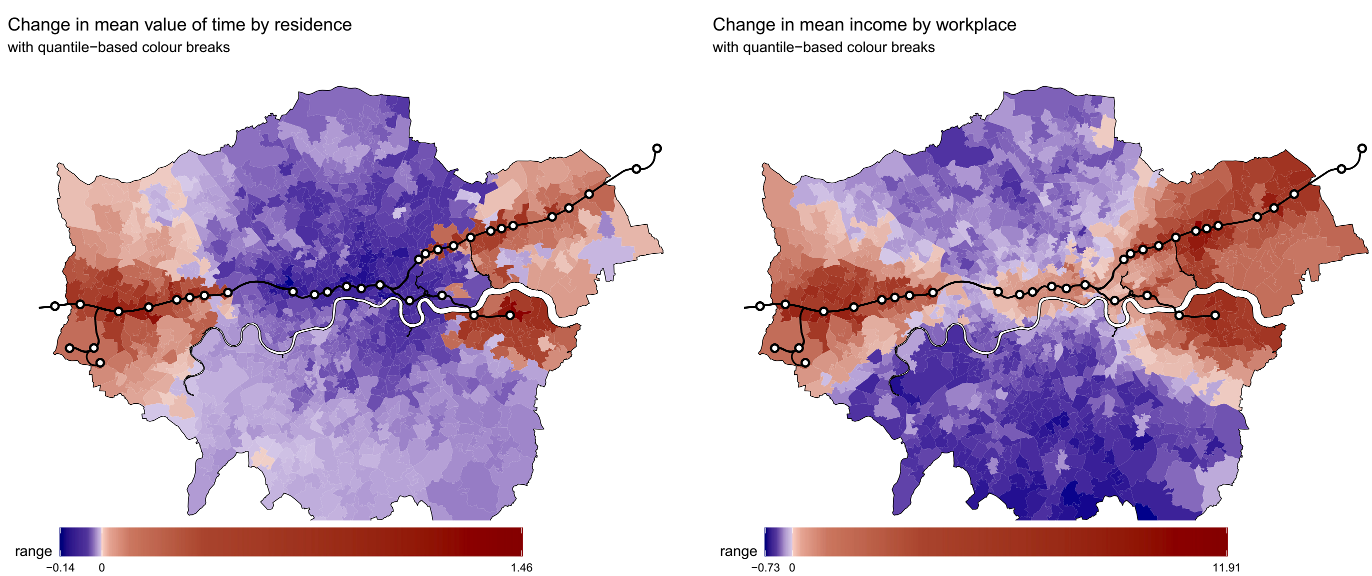

Fig. 11. Counterfactual outcomes: time valuation and labour supply.

Table 3
Benchmarking partial and general equilibrium appraisal methods.

Appraisal model	PE, no relocation		PE, with relocation		SGE
Value of time	homogeneous	heterogeneous	homogeneous	heterogeneous	n.a.
Direct user benefit	830.97***	752.84***	1082.4***	940.81***	810.04***
	(1.66)	(1.4)	(72.18)	(42.89)	(1.55)
Productivity gain	1766.59***	1766.59***	837.20	837.20	1342.07***
	(669.57)	(669.57)	(547.43)	(547.43)	(498.07)
MUoC dispersion					167.90***
					(59.15)
MUoH dispersion					−21.71**
					(10.84)
Land value uplift					34.61
					(53.11)
Fare revenues	1.33	1.33	1.33	1.33	1.33
	(26.27)	(26.27)	(26.27)	(26.27)	(26.27)
Total benefit	2598.88***	2520.75***	1920.93***	1779.34***	2334.23***
	(647.04)	(512.06)	(647.06)	(524.68)	(498.73)

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

The final column of Table 3 presents the welfare measure of the SGE model. In this case, the split between user benefits and productivity gains remains similar in magnitude to the PE model without relocation, although both quantities are somewhat lower. Although the productivity externality is also significant even at the 99% level, statistical uncertainty is considerably lower for the estimated direct user benefit. At the lower end of the 95% confidence band around the productivity gain, the estimate falls to less than half of the direct user benefit.

The remaining components of the aggregate welfare change include the welfare effects of the redistribution of consumption and housing patterns. The results indicate that consumption relocates towards places where the marginal utility of consumption is relatively high, generating a statistically significant welfare gain of £167,900 per day. By contrast, the relocation of housing consumption induces a mild but statistically significant welfare loss. The financial consequences of the project, including land value uplift and fare revenue generation, are small in magnitude and statistically uncertain; the confidence bands of these estimates do include zero.

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

Finally, we note that the welfare decomposition in Proposition 1 of Donald et al. (2025), and its adaptation in Eq. (30), are first order approximations that are best suited to modelling small deviations from the initial equilibrium. We also computed the alternative welfare measure of the SGE model using Eq. (31), which is not subject to approximation bias. This yielded an aggregate welfare change of £2,473.56 thousand per day (including landowner revenues), which is somewhat higher than our central estimate of £2,334.23. This approach has two drawbacks though. First, the aggregate welfare change cannot be decomposed into additive items. Second, solving for the compensating income requires a numerical optimisation step, which makes uncertainty quantification via bootstrapping computationally expensive. In our case, this cost was prohibitively high.

5.4. Sensitivity tests

We have tested the robustness of these findings in two alternative scenarios. The detailed results of these tests are provided in Appendix E. First, in Table E.1 we recalculate the appraisal results above with the UK-wide average agglomeration elasticity of $\eta = 0.044$ and a distance decay $\delta = -0.08$. Note that in this scenario we recalculate the vector of productivity fundamentals in line with Section 4.4, so that the pre-intervention spatial equilibrium is matched with the observed data through newly quantified vectors of productivity fundamentals. With this parameterisation, we find that the SGE welfare outcome remains of the same order of magnitude as the PE estimates, although it is somewhat higher than the PE result with no relocation. This confirms that it is unlikely that the divergence between the SGE and PE outcomes in the baseline model was driven by the relatively high agglomeration elasticity.

Second, in Table E.2 we check the sensitivity of our results with respect to the assumed magnitude of the relaxation of the floorspace construction restriction, i.e. $\bar{H}$. The PE model without relocation is unaffected by this assumption, as relocation is ignored in this model anyway. The PE model with relocation produces monotonically increasing welfare results. This finding is in line with intuition: as the transport improvement is combined with the possibility of developing more floorspace near the new stations, the employment densification unlocks further welfare gains and external agglomeration benefits. In fact, relaxing $\bar{H}$ in Central London is likely welfare enhancing in this model anyway, even in the absence of a complementary transport improvement. We observe an increasing pattern in the SGE results as well, but the rate of increase is even more pronounced. When $\bar{H}$ does not change at all around the Central London station, the welfare gain in SGE is lower than the PE result, likely due to job decentralisation and the associated loss of access to economic mass in the centre, which we observe when examining the spatial outcomes of the model with constant land use restrictions. By contrast, after a 5% relaxation of the floorspace constraint, the welfare gain is more than 50% higher than after a pure transport improvement with no land use intervention.

These results are in line with Tsivanidis (2026); Severen (2023) and Chen et al. (2024), which are QSM studies on urban transport improvements recognising the mutual dependence between transport and land-use policies, and the limited efficiency of the former when the latter remains unchanged. The combination of mass public transport development with urban planning measures, often referred to as Transit Oriented Development, has a large descriptive literature reviewed by Ibraeva et al. (2020). Thus, it is particularly remarkable that the current state-of-the-practice in transport appraisal is unable to handle the multiplier effect of land-use interventions in a coherent economic framework.

The high sensitivity of the SGE welfare results to local construction restrictions highlights the importance of making robust assumptions about the zoning interventions that are expected to accompany the transport improvement. However, inferring the appropriate magnitude of the change in $\bar{H}$ is not trivial in the case of the Elizabeth Line, even in an ex-post analysis. The reason is that in London and much of the United Kingdom, there are no explicit zoning regulations expressed in the form of height limits or maximum floorspace to land area ratios; construction rights are typically granted on a discretionary basis (Gurran and Bramley, 2017; Cheshire et al., 2018). Thus, it is difficult to know a priori what construction requests will be submitted and approved, or to identify ex-post which permissions were granted specifically because of the Elizabeth Line. There is substantial potential for further interdisciplinary research at the intersection of urban planning transport and real estate economics in connection with this matter. Until then, we believe that sensitivity testing is the right approach to quantify aggregate welfare effects under various scenarios.

To make further inferences about the external validity of the paper's core findings, in particular the magnitude of benefits in SGE relative to the mainstream PE appraisal, we performed a randomised experiment with a series of hypothetical transport improvements. Each Monte Carlo scenario assumes travel time reduction between a randomly selected central MSOA and a peripheral zone. The technical details of the experiment are documented in Appendix F, while the core results are shown in Fig. 12. We observe that across 250 randomly constructed counterfactual episodes, the high correlation between PE and SGE appraisal benefits remains remarkably robust. In relative terms, small improvements are more susceptible to diverging appraisal estimates. For example, when the SGE model predicts zero or even negative welfare effects, the PE method without relocation may still indicate non-negligible gains. The tendency of the PE approach with inelastic location choice to produce higher benefits also emerges as a consistent pattern.

5.5. Benchmarking against existing appraisal and evaluation reports

This methodological work does not intend to be a comprehensive cost-benefit analysis of the Elizabeth Line, but a comparison of our partial equilibrium results with official appraisal reports of the project may be useful for sense-checking our line of thinking. Buchanan and

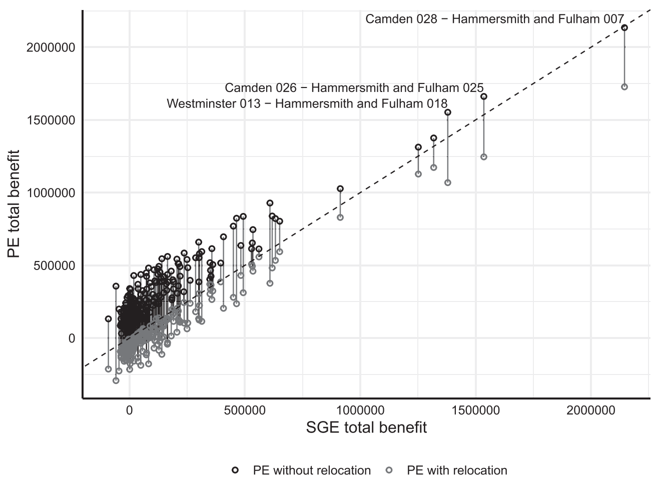

Fig. 12. Visual representation of randomly generated transport improvements. In each of the 250 steps of the randomised process, we selected a central and a peripheral MSOA and reduced travel times between the origin–destination pairs in their proximity. In each scenario, we computed the counterfactual spatial equilibrium and quantified the aggregate benefit according to the competing partial and general equilibrium methods. The dashed line is the 45° ray along which the PE results coincide with the SGE results.

Consulting (2007), one of the core appraisal studies that was completed just before construction began, reports that the conventional user benefit of the Elizabeth Line is £12.8 billion in 2002 present value terms for a 60-year period. They apply a 3.5% discount rate for the first forty years of the project and 3% in the remaining twenty years. Their calculation also assumes a 1.75% exogenous increase in productivity, which leads to a continuous increase in the value of time. Against this benchmark, our central PE estimate of direct user benefits is £809,000 per day. Assuming 253 workdays per year yields a yearly benefit of £200 million. The naive product of this figure and 60 years is £12.3 billion, very close to the official estimate. The 3.5 and 3% discount rates described above, combined with a constant yearly benefit of £0.2 billion, lead to a present value of £5.3 billion over 60 years. If we assume that this yearly benefit increases by 1.75% over the project duration, the discounted present value increases to £8.0 billion. Given that our model does not include the reduction in crowding in public transport, road congestion, and pollution externalities, these somewhat lower figures are in line with prior expectation.

The agglomeration benefits (part of the Wider Economic Benefits in the report) amount to 53.1 to 72.5% of the conventional user benefits in Buchanan and Consulting (2007). Our central estimates are significantly higher than these ratios, which is not surprising because our agglomeration elasticity is significantly higher than the DfT guidance from 2007. However, when we re-run the appraisal model with the UK-wide recommended 4.4% agglomeration elasticity in Table E.1 of Appendix E, the ratio of agglomeration benefits relative to direct user benefits drops to 61.5% and 54.4%, which are again close to what the official report provides. Therefore, we conclude that with comparable parameterisation, our results are of the same order of magnitude as the official appraisal of the Elizabeth Line from 2007.

We find it important to validate the predictive power of novel spatial models. In the case of some previous QSMs that rely on historic data spanning several decades or even a century (Ahlfeldt et al., 2015; Heblich et al., 2020; Koster et al., 2024), validation efforts have shown remarkable performance. However, in the case of the Elizabeth Line, validation is hindered by two important constraints. First, the railway line opened in May 2022, only three years ago, which is not sufficient to assume that the London economy has reached a new spatial equilibrium. Second, the line opened soon after the peak of the Covid-19 pandemic, which triggered transformative changes in commuting habits, remote work, and in particular the Central London property market. Thus, isolating the causal impact of the Elizabeth Line from the effects of Covid-19 is regarded in a recent ex-post evaluation of the project as a subject for future post-opening studies (Arup, 2024).

In general, we caution that even though our model combines three major markets (transport, labour, and housing), important urban mechanisms may still have been omitted. The model should not be considered a predominantly predictive tool that can be validated directly against observed data. We believe such models should always be paired with a causal ex post evaluation that can isolate, in data, the impact of the mechanisms included in the model.

6. Conclusions

This paper introduces a spatial general equilibrium model that combines the strengths of two influential branches of the literature: the empirical properties of quantitative spatial models and the theoretical foundations of computable spatial general equilibrium models for transport policy appraisal. The paper makes a contribution to the spatial economics literature by introducing a QSM featuring the leisure–labour trade-off with separate monetary and temporal budget constraints and a continuous decision margin on individual labour supply. This also enables us to derive an analytical formula for the marginal opportunity cost of time as a resource in monetary terms, that is, a spatially differentiated measure of the value of travel time savings. Taking a model of Greater London with nearly one thousand spatial units as a case study, we use this methodology to uncover the spatial distribution of travel time valuations, thereby providing important insights without a costly data collection effort.

Our study also contributes to the transport literature through a model that shares the theoretical properties of SCGE models, maintaining the desirable empirical properties of QSMs. The model is invertible, so that the contributions of exogenous geographic advantages and endogenous economic density on firm productivity (and urban amenities) can be disentangled and quantified for a large number of locations. The core structural parameters of the model are identified using structural estimating equations directly extracted from the theoretical model. We emphasise the importance of causality in the estimation of large-scale spatial equilibrium models. We argue that theoretically coherent causal parameter estimation reduces the uncertainty stemming from the ad-hoc parameter selection and associational (non-causal) calibration in previous transport SCGE models.

In a numerical application of the fully estimated model, we run a counterfactual simulation involving the hypothetical introduction of the Elizabeth Line in the census year of 2011, using historical data to represent the 2011 state of the London economy. The simulation predicts the reorganisation of economic activity in response to the transport policy. Then, in a second numerical exercise, we benchmark the welfare predictions of the QSM against traditional partial equilibrium transport appraisal methods.

One of the key messages of this paper is that quantitative spatial models, when suitably adapted for transport analysis, have the potential to become an established component of the transport appraisal toolbox. However, it is unlikely that a spatial equilibrium approach would entirely replace the mainstream CBA methodology. While spatial equilibrium analyses offer valuable insights into large-scale transformative interventions, applying them to smaller improvements with localised impacts introduces unnecessary complexity. Even for transformative projects, our preliminary numerical simulation of the Elizabeth Line in Greater London suggests that spatial equilibrium-based analyses produce aggregate welfare estimates similar in magnitude to those derived from an appropriately calibrated PE model and the common elasticity-based approximation of external agglomeration benefits.

The key advantage of spatial general equilibrium models is their ability to capture transport improvements' effects on *local* economic outcomes with high spatial granularity. These outcomes include the distributions of wages, employment, housing prices, and residential and workplace populations: metrics that policymakers and the public find more intuitive and accessible than abstract indicators such as the net present value of a welfare gain, the benefit-cost ratio, or the internal rate of return. By providing these spatially detailed but more comprehensible measures, SGE models enhance the transparency of cost-benefit analysis, ultimately fostering greater trust in transport appraisal within the public policy arena.

As briefly acknowledged at the beginning of Section 5, our model, in its current form, is not yet ready to perform a cost-benefit analysis in line with the official guidance manuals, or produce results that can be benchmarked against the official CBA reports for the Elizabeth Line. In pursuit of this goal, the paper unlocks a range of promising future research directions. One of the most pressing priorities is to enhance the representation of commuting behaviour by incorporating endogenous transport mode and route choice, alongside congestion and crowding externalities, following Tsivanidis (2026) and Allen and Arkolakis (2022), for example. Using data on mode-specific origin-destination matrices, incorporating mode choice is feasible in this paper's model variant as well, as demonstrated by Doffkay (2024). However, addressing route choice and congestion presents a more significant challenge, both empirically and computationally.

We identify three potential approaches to incorporating congestion into our model. (i) Anas and Liu (2007) integrate a standard stochastic traffic assignment process into their SCGE model. However, we find this approach challenging within our framework, as standard stochastic traffic assignment models lack a clear methodology for causally estimating the parameters of a link-level congestion cost function. Even if this issue were addressed, simulating counterfactual equilibria with one million origin-destination pairs or more and tens of thousands of route segments would present a significant computational challenge.[21] (ii) Under a unique set of assumptions, Allen and Arkolakis (2022) have recently addressed both the empirical and computational challenges of congestion modelling in a QSM. In their approach, delay on each road segment has a multiplicative impact on travel disutility, following the iceberg specification. Consequently, this method is incompatible with the framework of this paper and other models that rely on a leisure–labour trade-off. In our view, further work is required to validate these assumptions and prove the behavioural coherence of their model before it can be applied in a formal transport cost-benefit analysis. (iii) Koster (2024) endogenises the travel time matrix as a function of *effective traffic density* at the destination zone of commuting origin-destination pairs.[22] The concept assumes that congestion delay can be reasonably approximated by the volume of trips directed towards a commuter's workplace, as well as nearby workplaces, weighted by a distance decay term. Koster's empirical approach resembles the 'bathtub' (Arnott, 2013) and 'macroscopic fundamental diagram' (Daganzo and Geroliminis, 2008) models in the transport literature. However, his model overlooks route choice and the zones a commuter may travel through, indicating the need for further refinement to make it suitable for transport-oriented applications.

Another key step towards transforming our model into a comprehensive appraisal methodology is to expand its scope to encompass multiple trip purposes. At present, the model is limited to commuting flows. Future research will focus on developing microfoundations for both leisure and business travel. The empirical significance of non-commuting trips is substantial: commuting typically accounts for less than half of households' total transport consumption, a proportion that has further declined with the rise of remote work since the pandemic (see, e.g., Balbontin et al., 2024). Despite this, spatial models have rarely incorporated non-commuting trips. Among the exceptions, Anas and Liu (2007) and Lee and Tan (2024) account for shopping trips through endogenous consumption location choices, while Fajgelbaum et al. (2023) include leisure and business travel components in their long-distance rail model. However, neither of these models provides empirical methods to fully quantify their frameworks.

One significant challenge of integrating multiple trip purposes is that urban residents often engage in complex trip chains rather than simple, back-and-forth movements between home and a single activity location, as assumed in the current model. This complexity leads to a high-dimensional activity location choice problem. As an early effort to address this issue, Miyauchi et al. (2025) adapt the canonical quantitative urban model (incorporating travel time costs in an iceberg specification of utility) and tackle dimensionality with an importance sampling approach. These studies offer potential pathways for the future application of quantitative spatial models in transport appraisal.

CRediT authorship contribution statement

Daniel Hörcher: Conceptualization, Methodology, Software, Investigation, Visualization, Writing. **Daniel J. Graham:** Conceptualization, Methodology, Resources.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used ChatGPT (OpenAI), an online generative AI tool, in order to improve language quality. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Technical appendix of the analytical model

A.1. Household preferences

We begin the derivation of the analytical model with the household problem. First please note that the nested Cobb-Douglas specification in Eq. (1) can be rewritten in the form

$$U_{ij} = \gamma^{-\gamma}\left(\frac{L_{ij}}{1-\gamma}\right)^{1-\gamma}\left(\frac{C_{ij}}{\beta}\right)^{\gamma\beta}\left(\frac{H^R_{ij}}{1-\beta}\right)^{\gamma(1-\beta)} z_{ij}. \tag{A.1}$$

We provide the nested form in (1) to make the interpretation of expenditure shares β and γ easier. The full specification of the Lagrangian function associated with the maximisation of utility in (1) subject to constraints (2) and (3) is

$$\begin{aligned}\mathcal{L} = &\left(\frac{L_{ij}}{1-\gamma}\right)^{1-\gamma}\left[\frac{1}{\gamma}\left(\frac{C_{ij}}{\beta}\right)^{\beta}\left(\frac{H^R_{ij}}{1-\beta}\right)^{1-\beta}\right]^{\gamma} \\ &- \kappa_{ij}\left[p_i C_{ij} + q_i H^R_{ij} + x_{ij}\tau_{ij} - x_{ij}w_j\right] \\ &- \mu_{ij}\left[L_{ij} + x_{ij}(T_j + t_{ij}) - \bar{L}\right].\end{aligned} \tag{A.2}$$

First-order conditions are

$$\frac{\partial \mathcal{L}}{\partial x_{ij}} = -\kappa_{ij}(\tau_{ij} - w_j) - \mu_{ij}(T_j + t_{ij}) = 0; \tag{A.3}$$

$$\frac{\partial \mathcal{L}}{\partial C_{ij}} = \frac{\partial u_{ij}}{\partial C_{ij}} - \kappa_{ij} p_i = u_{ij}\gamma\left(\frac{C_{ij}}{\beta}\right)^{-1} - \kappa_{ij} p_i = 0; \tag{A.4}$$

[21] We believe that SGE models should not be expected to replace the role of standard four-stage traffic models commonly used in mesoscopic traffic simulation.

[22] This measure resembles the effective employment density terms commonly used in agglomeration research; see Eqs. (22) and (24) in this paper as an example.

$$\frac{\partial \mathcal{L}}{\partial H_{ij}^R} = \frac{\partial u_{ij}}{\partial H_{ij}^R} - \kappa_{ij} q_i = u_{ij}\gamma\left(\frac{H_{ij}^R}{1-\beta}\right)^{-1} - \kappa_{ij} q_i = 0; \tag{A.5}$$

$$\frac{\partial \mathcal{L}}{\partial L_{ij}} = u_{ij}\left(\frac{L_{ij}}{1-\gamma}\right)^{-1} - \mu_{ij} = 0. \tag{A.6}$$

Eqs. (4) and (5) follow directly from rearranging (A.3). After dividing (A.4) by (A.5), let us express the relationship between H_{ij}^R and C_{ij}

$$H_{ij}^R = \frac{(1-\beta)p_i}{\beta q_i} C_{ij} \tag{A.7}$$

Isolate the Lagrange multipliers using (A.4)–(A.6).

$$\begin{aligned} \kappa_{ij} &= u_{ij}\frac{\gamma\beta}{p_i C_{ij}} = u_{ij}\frac{\gamma(1-\beta)}{q_i H_{ij}^R} \\ \mu_{ij} &= u_{ij}(1-\gamma)L_{ij}^{-1} \end{aligned} \tag{A.8}$$

Substitute these expression into (5) and rearrange for C_{ij}.

$$C_{ij} = \frac{\gamma\beta}{1-\gamma}\frac{v_{ij}L_{ij}}{p_i} \tag{A.9}$$

Finally, express individual labour supply from the monetary budget constraint.

$$x_{ij} = \frac{p_i C_{ij} + q_i H_{ij}^R}{w_j - \tau_{ij}} \tag{A.10}$$

Based on these intermediate steps, we reach the utility-maximising leisure time, consumption and labour supply decisions as follows. Insert (A.10) into the time constraint in (3), substitute (A.7) and (A.9) above, and rearrange for L_{ij} to express the optimal leisure time:

$$L_{ij} = (1-\gamma)\bar{L}. \tag{A.11}$$

That is, γ is the share of the total time endowment that households allocate to working *and* commuting. As our functional form assumption does not impose any additional restrictions on the split between working time and commuting time, the intensive margin of the labour supply decision remains endogenous in this model.

Substitute (A.11) into (A.9) to get the demand for consumer good C_{ij}, and then use (A.7) to express the second part of (A.12), the optimal residential floorspace demand H_{ij}^R.

$$\begin{aligned} C_{ij} &= \beta\frac{\gamma\bar{L}v_{ij}}{p_i} \\ H_{ij}^R &= (1-\beta)\frac{\gamma\bar{L}v_{ij}}{q_i} \end{aligned} \tag{A.12}$$

Naturally, consumption decreases with the unit prices of goods and floorspace at the residential location. Next, v_{ij} enters this formula directly. That is, someone with a high value of time, or effective wage, is expected to consume more. The numerator in both equations, $\gamma\bar{L}v_{ij}$, is the daily net wage which depends on the cost and duration of commuting as well as the time allocated to work-related activities ($\gamma\bar{L}$). The indirect utility expressions in Eq. (7) of the main text follow from substituting (A.12) into the baseline specifications in (1).

As noted in the main text, the utility-maximising individual labour supply (6) follows directly by substituting (A.11) into the time constraint in (3). Equivalently, one can substitute C_{ij} and H_{ij}^R in (A.12) into (A.10) to derive individual labour supply x_{ij} in Eq. (6).

A.2. Production sector

Our derivation in this section follows the standard producer problem with a Cobb-Douglas production function. Based on the production function defined in (15), the cost minimisation problem has the following Lagrangian.

$$\mathcal{L} = w_j M_j^W + Q_j H_j^W - \lambda\left[A_j (M_j^W)^\alpha (H_j^W)^{1-\alpha} - Y_j\right] \tag{A.13}$$

The division of first-order conditions yields the regular equality between the marginal rate of substitution and the price ratio. After substituting this equation back into the production function, we get the following factor demand equations.

$$M_j^W = \left(\frac{\alpha}{1-\alpha}\frac{Q_j}{w_j}\right)^{1-\alpha}\frac{Y_j}{A_j}; \quad H_j^W = \left(\frac{1-\alpha}{\alpha}\frac{w_j}{Q_j}\right)^{\alpha}\frac{Y_j}{A_j} \tag{A.14}$$

The resulting cost function is

$$C_j(Y_j) = \frac{1}{1-\alpha}\left(\frac{1-\alpha}{\alpha}\right)^\alpha w_j^\alpha Q_j^{1-\alpha}\frac{Y_j}{A_j}. \tag{A.15}$$

Marginal revenue equals marginal cost at the profit-maximising output. Note that in this urban model we normalise the price of the homogeneous urban product to one, so that $p_i = 1\ \forall i$. Therefore,

$$1 = \frac{1}{1-\alpha}\left(\frac{1-\alpha}{\alpha}\right)^\alpha w_j^\alpha Q_j^{1-\alpha}\frac{1}{A_j}. \tag{A.16}$$

Rearranging this equation for Q_j and then w_j/Q_j, the profit maximising factor demand expressions in (A.14) become

$$\begin{aligned} M_j^W &= \left(\frac{\alpha A_j}{w_j}\right)^{\frac{1}{1-\alpha}} H_j^W; \\ H_j^W &= \left[\frac{(1-\alpha)A_j}{Q_j}\right]^{1/\alpha} M_j^W, \end{aligned} \tag{A.17}$$

or Eq. (14) in the main text. Finally, we consider that under the assumption of perfect competition and free entry to the market, profits drop to zero, so that

$$A_j (M_j^W)^\alpha (H_j^W)^{1-\alpha} - w_j M_j^W - Q_j H_j^W = 0. \tag{A.18}$$

Substitute (A.17) into the zero-profit constraint and rearrange for Q_j, the profit-maximising floorspace price under perfect competition, in Eq. (16) of the main text. Eq. (17) follows from (16) after a straightforward rearrangement.

A.3. Construction sector

The production function of the spatially differentiated construction sector in (18) is also Cobb-Douglas. The Lagrangian function corresponding to cost minimisation is

$$\mathcal{L} = Z_i + p_i^\ell \ell_i - \lambda\left[Z_i^{1-\psi}(\phi_i \ell_i)^\psi\right], \tag{A.19}$$

where we recall that capital (Z_i) is measured in the units of the homogeneous urban product, and thus its price is normalised to one, and ϕ_i is the multiplier in (19) that captures the tightness of the regulated local floorspace market. The ratio of the first-order conditions with respect to the input factors yield the following factor demand equations:

$$\begin{aligned} \ell_i &= \left(\frac{\psi-1}{\psi}\right)^{1-\psi}(p_i^\ell)^{\psi-1}\phi_i^{-\psi}H_i; \\ Z_i &= \left(\frac{1-\psi}{\psi}\right)^{\psi}(p_i^\ell)^{\psi}\phi_i^{-\psi}H_i. \end{aligned} \tag{A.20}$$

The cost function becomes

$$C_i(H_i) = \psi^{-\psi}\left(\frac{1}{1-\psi}\right)^{1-\psi}(p_i^\ell)^\psi\phi_i^{-\psi}H_i. \tag{A.21}$$

The unit price of floorspace supplied in location i is $\bar{q}_i$, the mean of local residential and commercial floorspace prices defined in Eq. (20). Thus, the first-order condition of profit maximisation in the construction sector is

$$\bar{q}_i = \psi^{-\psi}\left(\frac{1}{1-\psi}\right)^{1-\psi}(p_i^{\ell})^{\psi}\phi_i^{-\psi}. \tag{A.22}$$

Using the fact that ψ is the expenditure share spent on land, so that

$$p_i^{\ell} = \psi\frac{\bar{q}_i H_i}{\ell_i}, \tag{A.23}$$

(A.22) results in the following expression for the profit-maximising floorspace supply.

$$H_i = \phi_i(H_i)\left[(1-\psi)\,\bar{q}_i\right]^{(1-\psi)/\psi}\ell_i \tag{A.24}$$

Plugging the definition of $\phi_i(H_i)$ into this yields the final form of this expression in (21), which we use as an equilibrium condition.

Finally, the welfare calculations of this study in Section 5.3 require an explicit expression of land value, using endogenous outcomes in spatial equilibrium. Plugging the definition of $\bar{q}_i$ in (20) into (A.23), we derive

$$p_i^{\ell} = \psi\left(q_i H_i^R + q_i\xi_i H_i^W\right)\ell_i^{-1}. \tag{A.25}$$

This creates a relationship between floorspace prices and land value. Note that this comes directly from the zero profit constraint in the perfectly competitive floorspace sector: the total expenditure on land ($p_i^{\ell}\cdot\ell_i$) is the ψ fraction of the total revenue from floorspace rents.

Appendix B. Spatial equilibrium in counterfactual scenarios

The process of model quantification described in Section 4 ensures that the estimated structural parameters and residuals plugged into the equilibrium conditions reproduce the observed economic outcomes in our data. In other words, the state of the urban economy in which our data have been collected is an equilibrium of the model. We assess the impact of transport interventions in a comparative statics exercise. In other words, we modify some of the exogenous parameters, typically some of the elements of the transport time and monetary cost matrices, and compute the spatial equilibrium determined by the new parameter set.

In this section we provide the steps of the iterative process through which we compute the new equilibrium. Superscripts (0) and (1) denote the initial and the updated values of the variables that we endogenously update in each iteration, while $\hat{x}$ denotes the empirical estimate of any parameter x.

Step 1: Population distribution and labour supply. First we apply Eqs. (9) and (11) to compute the residential and workplace populations.

$$\begin{aligned}\lambda_{ij} &= \frac{\hat{X}_i\hat{E}_j\left[v_{ij}^{(0)}\left(q_i^{(0)}\right)^{\beta-1}\right]^{\gamma\hat{\epsilon}}}{\sum_r\sum_s\hat{X}_r\hat{E}_s\left[v_{rs}^{(0)}\left(q_r^{(0)}\right)^{\beta-1}\right]^{\gamma\hat{\epsilon}}}\\ N_i^R &= N\sum_j\lambda_{ij}\\ N_j^W &= N\sum_i\lambda_{ij}\end{aligned} \tag{B.1}$$

The counterfactual travel time matrix t_{ij} is substituted into Eq. (6) to compute x_{ij} for each origin–destination pair. Finally, Eq. (11) is used to aggregate the effective labour supply M_j^W.

Step 2: Local productivity and wages. Using the aggregate labour supply M_j^W derived above and the definition of local productivity in (22) and (24), we update the workplace-specific wage vector in line with Eq. (17).

$$\begin{aligned}A_j &= \hat{a}_j\left[\sum_s\exp(\hat{\delta}t_{sj})\,M_j^W\right]^{\hat{\eta}}\\ w_j^{(1)} &= \alpha A_j^{1/\alpha}\left(\frac{1-\alpha}{Q_j^{(0)}}\right)^{\frac{1-\alpha}{\alpha}}\end{aligned} \tag{B.2}$$

This enables us to update the v_{ij} matrix as well.

$$v_{ij}^{(1)} = \frac{w_j^{(1)}-\tau_{ij}}{T_j+t_{ij}} \tag{B.3}$$

Step 3: Floorspace prices. To determine the model's two floorspace price vectors, we first derive the aggregate residential floorspace demand for each location.

$$H_i^R = N_i^R\sum_j\lambda_{ij|i}\,H_{ij}^R, \tag{B.4}$$

which, after substituting (10) for $\lambda_{ij|i}$ and (A.12) for H_{ij}^R, yields

$$H_i^R = N_i^R(1-\beta)\frac{\gamma\bar{L}}{q_i^{(0)}}\sum_j\frac{\hat{E}_j\left(v_{ij}^{(1)}\right)^{\gamma\hat{\epsilon}}}{\sum_s\hat{E}_s\left(v_{is}^{(1)}\right)^{\gamma\hat{\epsilon}}}v_{ij}^{(1)}. \tag{B.5}$$

Aggregate commercial floorspace demand comes directly from (A.17),

$$H_i^W = \left[\frac{(1-\alpha)A_i}{Q_i^{(0)}}\right]^{1/\alpha}M_i^W, \tag{B.6}$$

in which we bring A_i and M_i^W from Steps 2 and 1, as detailed above. Total floorspace demand $H_i = H_i^R + H_i^W$ now enables us to update the residential floorspace price vector by substituting (20) into (21) and rearranging the latter for q_i.

$$q_i^{(1)} = \frac{1}{1-\psi}\left(H_i^R + H_i^W\hat{\xi}_i\right)^{-1}H_i^{\frac{1}{1-\psi}}\left[\left(1-H_i/\hat{\bar{H}}_i\right)\ell_i\right]^{\frac{\psi}{\psi-1}} \tag{B.7}$$

Finally, commercial floorspace prices are updated via $Q_i^{(1)} = \hat{\xi}_i\,q_i^{(1)}$.

Through the three steps above we update the seven vectors of location-specific variables, N_i^R, N_j^W, M_j^W, A_j, w_j, q_i and Q_j, ensuring that the equilibrium conditions are met in the labour, production and construction markets as well. Ahlfeldt et al. (2015) prove analytically that a structurally similar general equilibrium model, featuring Fréchet utility shocks in location choice and a multiplicative specification for utility, converges to a unique equilibrium.

Allen et al. (2024) analyse the equilibrium properties of a wide range of spatial interaction models using advances in contraction mapping theory. They characterise the conditions under which equilibrium exists and is unique using a matrix of uniform bounds on the elasticities of the spatial interaction functions with respect to local spatial outcomes, which measures the strength of such interactions. They derive sufficient and *globally necessary* conditions for the existence and uniqueness of equilibrium from the spectral radius of this matrix. In practical terms, in an urban model that follows the structure of spatial interactions in Ahlfeldt et al. (2015), uniqueness can be guaranteed when agglomeration forces are small relative to centrifugal forces. In this paper, we do not prove uniqueness analytically; however, our randomised numerical experiments did not reveal any evidence of equilibrium multiplicity.

In practice, multiple iterative algorithms can be used to achieve quick convergence. In this project we used a standard adaptive Method of Successive Averages algorithm to update the parameter vectors in $\theta = \{N_i^R, w_j, q_i\}$ as follows:

$$\theta^{(0)} := (1-\Phi)\cdot\theta^{(0)} + \Phi\cdot\theta^{(1)}, \tag{B.8}$$

where the updating parameter Φ is adaptive, i.e. it increases by 10% or decreases by 20% depending on whether the sum of the $|\theta^{(1)}-\theta^{(0)}|$ gaps is shrinking or widening between two consecutive iterations. Using this common approach and a standard PC, this model of nearly 1000 spatial units converges within around five minutes.

Appendix C. Empirical results on the effects of economic density

This section provides further details on the estimation of the generic parameters that capture the impact of economic density on local firm productivity and amenities.

C.1. Firm productivity

Our core results on firm productivity are summarised in Table 2 and Section 4.3. Fig. C.1 focuses on one specific outcome of the NLS model specified in (42) with control function (41); that is, model (4) in Table 2. It plots the $\hat{d}_r$ estimates and their 95% confidence bands. These coefficients quantify the contribution of the sequence of 2.5-minute travel time doughnuts to our measure of access to economic mass in an increasing order from location j, of which the TFP residual is the dependent variable in (42). We fix the value of the first travel time band at one, meaning that employment at location j itself is fully considered in the economic density measure.

The estimates in Fig. C.1 imply that the contribution of employment to access to economic mass is generally decreasing with travel time, confirming the effect of a distance decay. However, this decreasing pattern is not monotonic, and some of the coefficients of the nearest travel time bands are insignificant, which requires further attention. The latter result is not surprising. The low level of statistical significance is caused by the low number of observations at short distances: there are only very few pairs of MSOA centroids located within less than 10 min of travel time. The mildly upward sloping pattern between 15 and 30 min might indicate, counter-intuitively, that it matters more from a productivity perspective what firms in an MSOA can reach within 30 min than in 15 min. Based on the confidence bands, we cannot reject though the possibility that this pattern is flat or even very slightly downward sloping. These results remained robust in other specifications we tested during the project. Nevertheless, the pattern we observe here suggests that future research efforts may consider alternative distance decay specifications that differ from the negative exponential function used by the majority of the agglomeration literature.

In the main text we explain that our $\hat{\delta}$ estimate is based on the statistically significant $\hat{d}_r$ values only. The corresponding distance decay function is shown by the solid black line. The solid grey curve in Fig. C.1 is the distance decay function that we obtain by estimating δ including the non-significant coefficients as well. We find that the agglomeration externality decays more quickly in this case, but the pattern is not materially different.

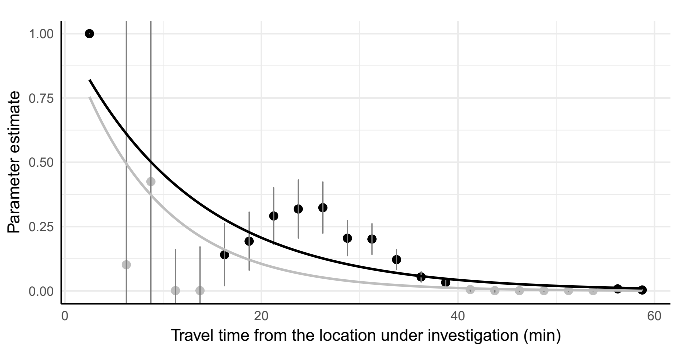

Fig. C.1. Estimates for travel-time bands in model (4) of Table 2. Black dots represent estimates that are significant at the 95% confidence level and the solid black line is the best-fitting decay function, i.e. $\exp(\hat{\delta}t_{rs})$, using significant estimates only. The grey dots and solid line correspond to non-significant estimates and the decay function based on the full sample.

Table C.1
Productivity elasticities in two subsamples based on distance from the CBD.

	Dependent variable: log productivity residual		
	(1)	(2)	(3)
Distance from CBD	Below median	*Full sample*	Above median
Mean employment density‡	4476	2763	1053
Density measure	Total emp.†	Total emp.†	Total emp.†
Method	NLS + CF	NLS + CF	NLS + CF
Productivity elasticity, η	0.175***	0.150***	0.081***
	(0.028)	(0.015)	(0.015)
Distance decay, δ	−0.090***	−0.080***	−0.054**
	(0.033)	(0.019)	(0.025)
Borough fixed effects	yes	yes	yes
RMSE	0.06	0.06	0.04
AIC	−1311.2	−2697.48	−1549.91
BIC	−1084.59	−2404.04	−1348.48
# of obs.	491	983	491

‡ : Measured in employees per km^2. Max value is 84,082 in the City of London.
† : Total employment is aggregated in 2.5-min travel time bands. Standard errors in parentheses, ***: 99%, **: 95%, *: 90%.

Table C.1 is a sensitivity test of our productivity model based on two sub-samples of the full dataset. Model (2) in this table is equivalent to Model (3) in Table 2.[23] The first column re-estimates η for the half of the MSOA sample located closer to the CBD. The table shows that the mean employment density is nearly twice as high as in the full sample; however, it is still just a fraction of the density of the City of London. We find that the agglomeration elasticity increases to 17.5% in this case. By contrast, when we consider more distant MSOAs closer to the periphery of the city, where the mean employment density is just a third of the entire city, the elasticity drops to 8.1%. Note that the distance decay parameter also seems to gradually decline as we move from the city centre towards the periphery, meaning that agglomeration spillovers spread over longer distances in the suburbs. This finding hints that (i) agglomeration economies are non-uniform, and (ii) as we move towards the city edge, the relatively high city-wide average agglomeration elasticity converges to the typical range of regional and nation-wide averages. The non-linear nature of agglomeration is in line with recent findings by Anupriya and Bansal (2023), but, to the best of our knowledge, it has not been tested in a QSM framework yet.

C.2. Residential and workplace amenities

The rest of this appendix section provides additional experimental results on the dependence of the local amenity residuals on the density of residential and employment activities. There is growing interest among researchers and practitioners in understanding agglomeration (more generally, economic density) effects beyond firm productivity (see e.g., Ahlfeldt and Pietrostefani, 2019) and in incorporating such effects in transport appraisal. For example, better connectivity may have external effects on amenities and the attractiveness of urban locations, which is a market failure currently ignored in the official CBA guidelines. From a methodological perspective, replacing the TFP residual ($\hat{A}_j$) with residential or workplace amenity residuals ($\hat{X}_i$ and $\hat{E}_j$) in the empirical models of Section 4.3 is an appealing possibility.

In the models reported in Table C.2, the dependent variable is $\log \hat{X}_i$, that is, the logarithm of the residential amenity residual recovered via model inversion (see Section 4.2). We find it natural to use population-related mass measures to explain residential amenities. In models (1) and (2) our choice is the population density, i.e. the ratio of the nighttime population and the area of each MSOA. In models (3) and (4) we broaden this definition to the population of all MSOAs

[23] We use Model (3) in this case because its computational time is significantly shorter than Model (4)'s while the resulting η and δ estimates are similar.

Table C.2
The elasticity of residential amenities with respect to population density.

	Dependent variable: log residential amenity			
Mass measure	(1) Population density	(2)	(3) Population in 10 min	(4)
Method	2SLS	CF	2SLS	CF
log mass measure	0.280***	0.280***	1.027***	1.027***
	(0.103)	(0.099)	(0.119)	(0.119)
Borough fixed effects	yes	yes	yes	yes
RMSE	0.58	0.56	0.60	0.55
AIC	−1018.48	1726.95	−940.88	1692.46
BIC	−852.2	1903.01	−774.6	1868.52
# of obs.	983	983	983	983

†: Total employment is aggregated in 2.5-min travel time bands.
Standard errors in parentheses, ***: 99%, **: 95%, *: 90%.

Table C.3
The elasticity of workplace amenities with respect to population density.

	Dependent variable: log workplace amenity			
Mass measure	(1) Employment density	(2)	(3) Employment in 10 min	(4)
Method	2SLS	CF	2SLS	CF
log mass measure	−0.200**	−0.187**	−0.029	−0.029
	(0.09)	(0.09)	(0.242)	(0.242)
Borough fixed effects	yes	yes	yes	yes
RMSE	1.00	0.62	0.91	0.57
AIC	64.72	1915.39	−124.67	1739.88
BIC	231	2091.45	41.61	1915.94
# of obs.	983	983	983	983

†: Total employment is aggregated in 2.5-min travel time bands.
Standard errors in parentheses, ***: 99%, **: 95%, *: 90%.

within a travel time of 10 min. Unobserved local characteristics and reverse causality raise endogeneity concerns in the same way as in the productivity models. Thus, in models (1) and (3) we instrument the mass measure by the 1841, 1861, 1881, 1901 and 1921 population density and a third-order polynomial of the distance from the CBD. In models (2) and (4) we use the control function approach introduced in Section 4.3, with the same historical and geographical instruments.

Higher population density may affect amenities in various ways: for example, density implies that more services may be available locally and there is a higher chance to interact within the local community. At the same time, population density increases the chance of friction between inhabitants, in the form of congestion, pollution and traffic externalities, for example. Thus, we do not have prior expectations about the sign of the estimated elasticities. Table C.2 reports statistically significant, positive, and relatively high elasticities. In general, the elasticity is higher when the total population within ten minutes is used as the mass measure. However, the estimates display substantial sensitivity with respect to the dependent variable: a 10% increase in population density is associated with a 2.8% increase in local amenities, whereas the population within ten minutes exhibits an almost one-to-one relationship with the residential amenity residual.

Table C.3 contains estimation results for the workplace amenities. The dependent variable is replaced with $\log \hat{E}_j$, the logarithm of the workplace amenity residual, and the mass measure is now the employment density (models 1 and 2) and total employment within 10 min (models 3 and 4). The latter pair of models did not yield statistically significant estimates. The estimates in models (1) and (2) with employment density are both significant at the 90% level. The elasticities indicate that workplace amenities are negatively affected by employment densities: this result hints that congestion and other nuisance factors are important determinants of how attractive a workplace location is. This result is particularly important due to its sign: in an evaluation context, negative workplace amenity externalities may at least partly neutralise the positive productivity externalities that the literature has devoted more attention to.

The experimental results reported in Tables C.2 and C.3 suggest that access to nearby residential and workplace population masses does have a considerable causal effect on residential and workplace amenities. Moreover, the magnitude of these effects is generally high, i.e. the estimated amenity elasticities are comparable to, or even higher than, the productivity elasticities in the main text. However, the coefficients and even their sign and statistical significance are highly dependent on the choice of mass measure. The literature on amenities and agglomeration is still in its early stages and provides limited guidance for selecting a preferred model among those discussed above. Given the high variability in the estimated coefficients and their significance, we have decided not to incorporate endogenous amenity externalities into the appraisal exercise presented in Section 5.3. This remains a highly relevant subject for future research.

Appendix D. Welfare decomposition in spatial equilibrium

This section provides the full derivation of the additive welfare decomposition introduced in Section 3.5.2 of the main text. Our discussion is a summary of the welfare decomposition in Donald et al. (2025) with the notation and model structure adapted to our urban model.

The derivation begins with two lemmas from Hofbauer and Sandholm (2002). The first states that the equilibrium residence–workplace choices can be represented as the solution of a problem that maximises an additive objective function; the sum of utilities under local consumption levels and a function of the vector of location choice probabilities.

$$\mathbb{E}[U_{ij}] = \max_{\{\lambda_{ij}\}} \sum_{ij} \lambda_{ij} \cdot u_{ij}(L_{ij}, C_{ij}, H^R_{ij}) - \Psi(\{\lambda_{ij}\}) \quad \text{subject to} \quad \sum_{ij} \lambda_{ij} = 1 \tag{D.1}$$

The solution of this problem coincides with expected utility in the location choice situation: $\mathbb{E}[U_{ij}] = \mathbb{E}[\max_{ij}\{u_{ij}(\cdot) \cdot z_{ij}\}]$. The second lemma states that any decentralised equilibrium allocation represented by matrix $\boldsymbol{X}$ of endogenous location-specific variables solves a pseudo-planning problem specified as

$$\begin{aligned}
W = & \max_{\{\mathbb{E}[U_{ij}], \boldsymbol{X}\}} \mathcal{W}\left(\mathbb{E}[U_{ij}]\right), \\
& \text{subject to} \\
& \mathbb{E}[U_{ij}] = \sum_{ij} \lambda_{ij} \cdot u_{ij}(L_{ij}, C_{ij}, H^R_{ij}) - \Psi(\{\lambda_{ij}\}), \\
& \{\lambda_{ij}\} \in \arg \max_{\{\tilde{\lambda}_{ij}\}: \sum_{ij} \tilde{\lambda}_{ij} = 1} \sum_{ij} \tilde{\lambda}_{ij} \cdot u_{ij}(L_{ij}, C_{ij}, H^R_{ij}) - \Psi(\{\tilde{\lambda}_{ij}\}), \\
& \sum_j A_j f^y_j(M^W_j, H^W_j) = \sum_{ij} N \lambda_{ij}(\boldsymbol{L}_{ij}, \boldsymbol{C}_{ij}, \boldsymbol{H}^R_{ij}) \cdot C_{ij}, \\
& \sum_i N \lambda_{ij}(\boldsymbol{L}_{ij}, \boldsymbol{C}_{ij}, \boldsymbol{H}^R_{ij}) \cdot x_{ij} = M^W_j \quad \forall j, \\
& f^h(Z_j, \ell_j) = \left(\sum_k N \lambda_{jk}(\boldsymbol{L}_{jk}, \boldsymbol{C}_{jk}, \boldsymbol{H}^R_{jk}) \cdot H^R_{jk} \right) - H^W_j \quad \forall j, \\
& \mathbb{L}_j = \ell_j \quad \forall j.
\end{aligned} \tag{D.2}$$

In this problem, $\mathcal{W}(\cdot)$ transforms household utility into social welfare. As in Donald et al. (2025), $\mathcal{W}(\cdot)$ could include multiple types representing income groups or other segmentations of society, and type-specific Λ values could be used as welfare weights. Due to constraints in data availability, in this paper we restrict the analysis to one representative household, and so the purpose of $\mathcal{W}(\cdot)$ is limited to turning expected household utility into money-metric welfare through the normalisation introduced in Eq. (29). The first two constraints in (D.2) reflect utility maximising location choices according to (D.1). The remaining four sets of constraints ensure market clearing in the remaining markets: equalising the quantities supplied and demanded in the local production, labour, construction and land markets, respectively, according to our derivations in Sections 3 and Appendix A.

We extend the objective function in (D.2) with two terms: fare revenues for the public transport operator and the income of absentee landlords. Thus, the Lagrangian of the problem becomes

$$\begin{aligned}\mathcal{L} = \mathcal{W}\Bigg(&\sum_{ij} N\lambda_{ij}(\boldsymbol{L}_{ij}, \boldsymbol{C}_{ij}, \boldsymbol{H}^R_{ij}) \cdot u_{ij}(L_{ij}, C_{ij}, H^R_{ij}) - \Psi\left(\{\lambda_{ij}(\boldsymbol{L}_{ij}, \boldsymbol{C}_{ij}, \boldsymbol{H}^R_{ij})\}\right)\Bigg) \\ &+ \sum_{ij} N\lambda_{ij}(\boldsymbol{L}_{ij}, \boldsymbol{C}_{ij}, \boldsymbol{H}^R_{ij}) \cdot x_{ij}\tau_{ij} + \sum_j p^\ell_j \cdot \ell_j \\ &+ \tilde{p}\left[\sum_j A_j f^y_j(M^W_j, H^W_j) - \sum_{ij} N\lambda_{ij}(\boldsymbol{L}_{ij}, \boldsymbol{C}_{ij}, \boldsymbol{H}^R_{ij}) \cdot C_{ij}\right] \\ &+ \sum_{ij} \tilde{w}_{ij}\left[N\lambda_{ij}(\boldsymbol{L}_{ij}, \boldsymbol{C}_{ij}, \boldsymbol{H}^R_{ij}) \cdot x_{ij} - M^W_{ij}\right] \\ &+ \sum_j \tilde{q}_j\left[f^h(Z_j, \ell_j) - \left(\sum_k N\lambda_{jk}(\boldsymbol{L}_{jk}, \boldsymbol{C}_{jk}, \boldsymbol{H}^R_{jk}) \cdot H^R_{jk}\right) - H^W_j\right] \\ &+ \sum_j \tilde{p}^\ell_j\left[\mathbb{L}_j - \ell_j\right],\end{aligned} \tag{D.3}$$

which is equivalent to Eq. (28) in the main text with the notational simplification $\Theta = (\{L_{ij}\}, \{C_{ij}\}, \{H^R_{ij}\})$. Donald et al. (2025) show that in equilibrium, the Lagrange multipliers of the market clearing conditions equal the relevant market prices, so that $\tilde{p} = p$, and $\tilde{w}_{ij} = w_j - \tau_{ij}$, $\tilde{q}_j = q_j$ and $\tilde{p}^\ell_j = p^\ell_j$ for all i and j.

Let us consider the welfare impact of a small change in travel times on origin–destination pair kl. Given the above-mentioned property of the Lagrange multipliers, the definition in Eq. (29), and the application of the envelope theorem to the pseudo-planning problem, the associated welfare change is

$$\begin{aligned}\frac{dW}{dt_{kl}} = &\sum_{ij}\left[N_{ij}\Lambda\frac{\partial u_{ij}}{\partial L_{ij}}\frac{dL_{ij}}{dt_{kl}} + N_{ij}\left(\Lambda\frac{\partial u_{ij}}{\partial C_{ij}} - p\right)\frac{dC_{ij}}{dt_{kl}} + N_{ij}\left(\Lambda\frac{\partial u_{ij}}{\partial H^R_{ij}} - q_i\right)\frac{dH^R_{ij}}{dt_{kl}}\right] \\ &+ \sum_{ij}\tau_{ij}\frac{d(N_{ij}x_{ij})}{dt_{kl}} + \sum_j \ell_j\frac{dp^\ell_j}{dt_{kl}} \\ &+ \sum_{ij}\left((w_j - \tau_{ij})x_{ij} - pC_{ij} - q_i H^R_{ij}\right)\frac{dN_{ij}}{dt_{kl}} \\ &+ \sum_j p f^y_j \frac{dA_j}{d\rho_j}\frac{d\rho_j}{dt_{kl}}.\end{aligned} \tag{D.4}$$

Note that the marginal utility of leisure time is $\partial u_{ij}/\partial L_{ij} = \mu_{ij}$, and, from the leisure time constraint in (3), $L_{ij} = \bar{L} - x_{ij}(T + t_{ij})$, so that $dL_{ij}/dt_{ij} = -x_{ij}$. Similarly, recall that the marginal utility of income is κ_{ij}, from which $\partial u/\partial C_{ij} = p\kappa_{ij}$ and $\partial u/\partial H^R_{ij} = q_i\kappa_{ij}$. Recall the monetary budget contraint in (2), which implies that the third line in (D.4) is zero. Finally, given our specification of local productivity in (22), $dA_j/d\rho_j = \eta A_j/\rho_j$. We apply these consequences of our model definitions to rewrite Eq. (D.4) in a total differential form for a transport project that affects travel times on multiple origin–destination pairs in the network.

$$\begin{aligned}dW = &\underbrace{\sum_{ij}\Lambda\mu_{ij}N_{ij}x_{ij} \cdot (-dt_{ij})}_{(i)} + \underbrace{\sum_{ij}N_{ij}(\Lambda\kappa_{ij} - 1)\,dC_{ij}}_{(ii)} + \underbrace{\sum_{ij}N_{ij}(\Lambda\kappa_{ij} - 1)\,dH^R_{ij}}_{(iii)} \\ &+ \underbrace{\sum_{ij}\tau_{ij} \cdot d(N_{ij}x_{ij})}_{(iv)} + \underbrace{\sum_j \ell_j \cdot dp^\ell_j}_{(v)} + \underbrace{\sum_j \frac{p\eta Y_j}{\rho_j}d\rho_j}_{(vi)}\end{aligned} \tag{D.5}$$

To further simplify items (ii) and (iii) in this expression, let us consider the following covariance rules.

$$\begin{aligned}&\mathrm{Cov}[X, Y] = \mathbb{E}[X \cdot Y] - \mathbb{E}[X] \cdot \mathbb{E}[Y] \\ &\mathbb{E}[X] = \sum_{ij}\lambda_{ij}X_{ij}\end{aligned} \tag{D.6}$$

Thus, we rewrite items (ii) in (D.5) as

$$\begin{aligned}&\sum_{ij}N_{ij}(\Lambda\kappa_{ij} - 1)\,dC_{ij} = N\sum_{ij}\lambda_{ij}\left(\Lambda - \frac{1}{\kappa_{ij}}\right)\kappa_{ij}\,dC_{ij} \\ &= N \cdot \{\mathrm{Cov}[\Lambda - \kappa^{-1}_{ij}, \kappa_{ij}\,dC_{ij}] + \mathbb{E}[\Lambda - \kappa^{-1}_{ij}] \cdot \mathbb{E}[\kappa_{ij}\,dC_{ij}]\} \\ &= N \cdot \mathrm{Cov}[\Lambda - \kappa^{-1}_{ij}, \kappa_{ij}\,dC_{ij}]\end{aligned} \tag{D.7}$$

In the final step we considered that Λ is independent from locations, so that $\mathbb{E}[\Lambda - \kappa^{-1}_{ij}] = \Lambda - \mathbb{E}[\kappa^{-1}_{ij}]$, which equals zero under the normalisation introduced in (29). We rewrite item (iii) in (D.5) through the same steps to get

$$\sum_{ij}N_{ij}(\Lambda\kappa_{ij} - 1)\,dH^R_{ij} = N \cdot \mathrm{Cov}[\Lambda - \kappa^{-1}_{ij}, q_i\kappa_{ij}\,dH^R_{ij}]. \tag{D.8}$$

By plugging (D.6) and (D.8) into (D.5), we reach the additive decomposition of the total welfare effect in Eq. (30) in the main text. This result adapts Proposition 1 in Donald et al. (2025) to the present paper's model and notation.

Appendix E. Additional results on the Elizabeth line appraisal

Fig. E.1 breaks down the aggregate wider economic impact measures in Section 5.3 by plotting the external welfare change in Eq. (27) for each MSOA separately. Specifically, it plots the agglomeration benefit we get in the partial equilibrium model with relocation versus without relocation. We also add a 45 degree line corresponding to identical values on the two axes to reveal whether relocation adds (above the line) or reduces (below the line) local productivity. We find a mixed pattern. We observe a cluster of locations that realise relatively high gains without relocation and then their benefits further increase when location choice becomes endogenous—this is what common-sense intuition would anticipate. However, there is another group of outliers that showcase reasonably high gains in the former case but then the impact of agglomeration on welfare turns into even higher negative values with relocation. This showcases that modelling the spatial reorganisation of economic activity is crucial in predicting how a transport scheme is expected to affect local productivity. Limiting the appraisal exercise to transport cost reductions in a static environment may seriously misinform policy-makers about the presence of losers and the spatial distribution of those who gain and lose from place-based policies.

Tables E.1 and E.2 document the outcomes of two sensitivity tests for the main appraisal results reported in Table 3. In Table E.1 we recalculate all appraisal results by assuming an agglomeration elasticity

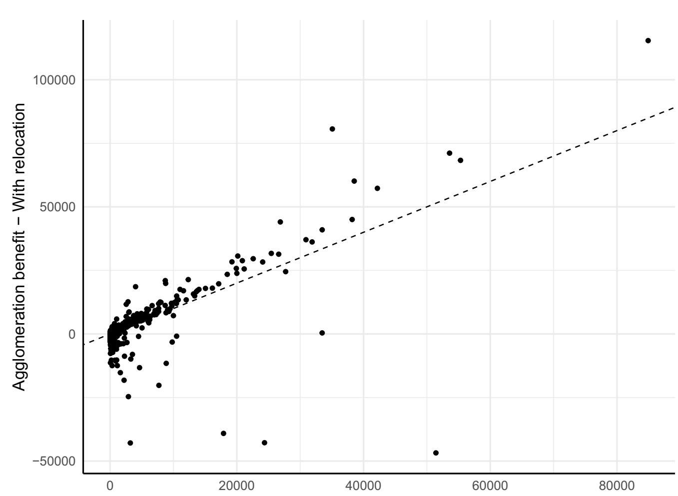

Fig. E.1. Agglomeration (productivity) benefits with and without relocation, plotted separately for each MSOA of Greater London. The dashed line connects equal values on the two axes.

Table E.1
Sensitivity test: Fig. 3 with lower agglomeration elasticity: $\eta = 0.044$, $\delta = -0.08$.

Appraisal model	Static partial equilibrium		Dynamic partial equilbrium		SGE
Value of time	homogeneous	heterogeneous	homogeneous	heterogeneous	n.a.
Direct user benefit	831	753	1026	918	810
Productivity gain	511	511	533	533	641
MUoC dispersion					74
MUoH dispersion					−7
Land value uplift					114
Fare revenues	42	42	42	42	42
Total benefit	1384	1306	1601	1493	1673

Units: Surplus in thousands of GBP per day.

Table E.2
Sensitivity test: Total benefit in Fig. 3 with different exogenous zoning (floorspace restriction) policies.

Appraisal model	Static partial equilibrium		Dynamic partial equilbrium		SGE
Value of time	homogeneous	heterogeneous	homogeneous	heterogeneous	n.a.
Restriction unchanged	2599	2521	1597	1454	1912
2% relaxation (baseline)	2599	2521	1921	1779	2334
5% relaxation	2599	2521	2376	2235	2926

Units: Surplus in thousands of GBP per day.

of $\eta = 0.044$ and a distance decay of $\delta = -0.08$. These parameters are often considered as a median of the existing estimates in the literature (see Graham and Gibbons, 2019) as well as approximations of the parameters recommended by the UK Transport Analysis Guidance for the entire country.

Table E.2 is a sensitivity test in which we vary the extent to which parameter $\bar{H}_i$, the exogenous floorspace construction constraint in Eq. (19), is relaxed in the counterfactual scenario of the Elizabeth Line. In this table we include the total benefit only (i.e., the final row of Table 3 in the main text). The results of these sensitivity tests are interpreted in Section 5.4.

Appendix F. Randomised transport improvements

The core findings of our simplified appraisal of the Elizabeth Line in Section 5 are that (i) the SGE appraisal results are of the same order of magnitude as those from the mainstream PE appraisal, and (ii) the SGE benefits are somewhat lower than the PE benefits when we ignore the relocation of households and firms in travel demand and firm productivity estimation, and somewhat higher when we take such relocation into account. These findings are specific to the Elizabeth Line, which limits their external validity and applicability to other policy interventions.

To be able to say more about external validity, we have designed a randomised simulation experiment. In a repeated process, we randomly selected a central MSOA within four kilometres of the City of London (the central business district) and a more peripheral MSOA more than eight kilometres from the CBD. For each randomly selected origin-destination pair, we assumed that a hypothetical transport service is established that significantly reduces the travel time between them. To extend the number of residence-workplace pairs affected by the synthetic policy, we drew a four kilometre catchment area around the hypothetical stations, and allowed travel times to decrease between every origin-destination pair that connects the two catchments. More specifically, if the hypothetical stations are at the centroids of MSOAs i and j, and locations l and k are within the four kilometre ranges of i and j, then we multiplied the travel time t_{kl} by

$$0.5 \cdot \left(1 + \frac{d(k,i) + d(l,j)}{8\,000}\right), \tag{F.1}$$

where $d(x_1, x_2)$ is the Euclidean distance between centroids x_1 and x_2 in metres. This implies that travel times are halved between the MSOAs directly connected by the hypothetical transport service, while the travel time between two centroids lying on the four kilometre circles remains unaffected. Fig. F.1 illustrates this process for one randomly selected origin-destination pair. The colour gradients within the catchment areas are proportional to the expected reduction in travel time.

This process has been repeated 250 times. In each hypothetical scenario, we computed the counterfactual spatial equilibrium and quantified the appraisal metrics defined in Section 3.5 and evaluated in Table 3 for the Elizabeth Line. The results of this experiment are visualised in Fig. 12 and discussed in Section 5.4.

The results of this randomised experiment allow us to strengthen the external validity of some of the paper's core findings. However, the Monte Carlo experiment also has its own specific features and limitations. For example, one may argue that transport infrastructure is typically linear and connects more than two distinct locations in

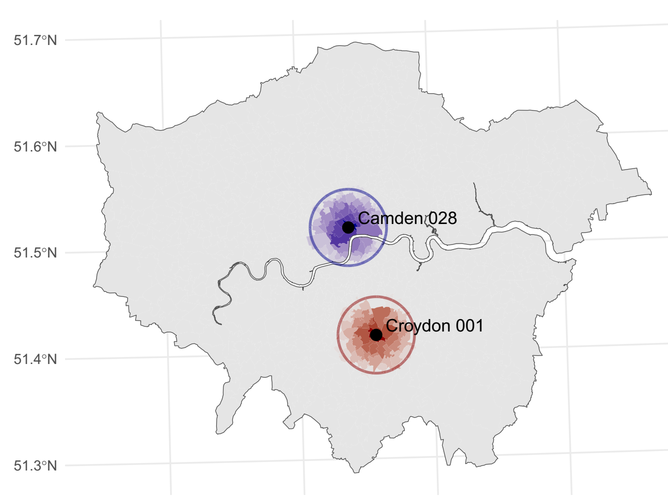

Fig. F.1. One example of a randomly selected origin–destination pair. Travel times are assumed to decrease by up to 50%, in proportion to the distance from the hypothetical new station (as shown by the colour grading on the map).

the city. The present experiment has been designed to ensure simple implementation and transparency. Future research may repeat the exercise with more realistic randomly generated transport corridors. Second, as described above, we combined a central and a peripheral location in each round of the simulation. Although this rule slightly restricted the randomness of the experiment, earlier attempts with fully random origin–destination pairs often selected very remote locations with virtually zero traffic before the intervention, which made the simulations less meaningful. Finally, one may note that the present results are restricted to London because they rely on the baseline data and parameters we estimated. Naturally, real-world or randomised appraisal exercises may lead to different results in other geographical contexts.

Data availability

Data will be made available on request.

References

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

Ahlfeldt, G.M., Pietrostefani, E., 2019. The economic effects of density: a synthesis. J. Urban Econ. 111, 93–107.

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

Ahlfeldt, G.M., Carozzi, F., Makovsky, L., 2023. A micro-geographic house price index for England and Wales. Cep occasional paper. The London School of Economics and Political Science.

Allen, T., Arkolakis, C., 2022. The welfare effects of transportation infrastructure improvements. Rev. Econ. Stud. 89 (6), 2911–2957.

Allen, T., Arkolakis, C., Li, X., 2024. On the equilibrium properties of spatial models. Am. Econ. Rev. Insights 6 (4), 472–489.

Almagro, M., Barbieri, F., Castillo, J.C., Hickok, N.G., Salz, T., 2024. Optimal urban transportation policy: evidence from Chicago. NBER Working Paper No. w32185..

Anas, A., 2020. The cost of congestion and the benefits of congestion pricing: a general equilibrium analysis. Transp. Res. B Methodol. 136, 110–137.

Anas, A., Chang, H., 2023. Productivity benefits of urban transportation megaprojects: a general equilibrium analysis of <<grand Paris express>>. Transp. Res. B Methodol. 174, 102746.

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

Anas, A., Liu, Y., 2007. A regional economy, land use, and transportation model (RELU-TRAN©): formulation, algorithm design, and testing. J. Reg. Sci. 47 (3), 415–455.

Anupriya, G., D.J, Bansal, P., 2023. Quantification of non-linear effects in agglomeration economies for transport appraisals. In: Paper Presented at the 16th World Conference on Transport Research, 2023, Montreal, Canada.

Arnott, R., 2007. Congestion tolling with agglomeration externalities. J. Urban Econ. 2 (62), 187–203.

Arnott, R., 2013. A bathtub model of Downtown traffic congestion. J. Urban Econ. 76, 110–121.

Arup, 2024. Elizabeth Line Post-Opening Evaluation. Transport Outcomes: Interim Findings Report'. Report Comissioned by Transport for London and the UK Department for Transport (20 June 2024).

Balbontin, C., Hensher, D.A., Beck, M.J., 2024. The influence of working from home and underlying attitudes on the number of commuting and non-commuting trips by workers during 2020 and 2021 pre-and post-lockdown in Australia. Transp. Res. a Policy Pract. 179, 103937.

Bansal, P., Hörcher, D., Graham, D.J., 2022. A dynamic choice model to estimate the user cost of crowding with large-scale transit data. J. R. Stat. Soc. a Stat. Soc. 185 (2), 615–639.

Barwick, P.J., Li, S., Waxman, A., Wu, J., Xia, T., 2024. Efficiency and equity impacts of urban transportation policies with equilibrium sorting. Am. Econ. Rev. 114 (10), 3161–3205.

Becker, G.S., 1965. A theory of the allocation of time. Econ. J. 493–517.

Bento, A., Roth, K., Waxman, A., 2024. The value of urgency: evidence from real-time congestion pricing. J. Polit. Econ. Microecon. 2 (4), 786–851.

Börjesson, M., Eliasson, J., 2019. Should values of time be differentiated? Transp. Rev. 39 (3), 357–375.

Buchanan, C., Consulting, V., 2007. The Economic Benefits of Crossrail: Final Report. Final Report of Project No 125721 (May 2007).

Buchholz, N., Doval, L., Kastl, J., Matejka, F., Salz, T., 2025. Personalized pricing and the value of time: evidence from auctioned CAB rides. Econometrica 93 (3), 929–958.

Cadot, O., Röller, L.-H., Stephan, A., 2006. Contribution to productivity or pork barrel? The two faces of infrastructure investment. J. Public Econ. 90 (6–7), 1133–1153.

Chen, L., Hasan, R., Jiang, Y., Parkhomenko, A., 2024. Faster, taller, better: transit improvements and land use policies. J. Dev. Econ. 171, 103322.

Cheshire, P., Hilber, C.A., Koster, H.R., 2018. Empty homes, longer commutes: the unintended consequences of more restrictive local planning. J. Public Econ. 158, 126–151.

Combes, P.-P., Gobillon, L., 2015. , The empirics of agglomeration economies. In: Handbook of Regional and Urban Economics, vol. 5. Elsevier, pp. 247–348.

Combes, P.-P., Duranton, G., Gobillon, L., 2019. The costs of agglomeration: house and land prices in French cities. Rev. Econ. Stud. 86 (4), 1556–1589.

Daganzo, C.F., Geroliminis, N., 2008. An analytical approximation for the macroscopic fundamental diagram of urban traffic. Transp. Res. B Methodol. 42 (9), 771–781.

Davis, M.A., Ortalo-Magné, F., 2011. Household expenditures, wages, rents. Rev. Econ. Dyn. 14 (2), 248–261.

De Borger, B., Van Dender, K., 2003. Transport tax reform, commuting, and endogenous values of time. J. Urban Econ. 53 (3), 510–530.

De Borger, B., Wuyts, B., 2011. The structure of the labor market, telecommuting, and optimal peak period congestion tolls: a numerical optimization model. Reg. Sci. Urban Econ. 41 (5), 426–438.

Delventhal, M., Parkhomenko, A., 2024. Spatial implications of telecommuting, Available at SSRN 3746555.

Dericks, G.H., Koster, H.R., 2021. The billion pound drop: the blitz and agglomeration economies in London. J. Econ. Geogr. 21 (6), 869–897.

DeSerpa, A.C., 1971. A theory of the economics of time. Econ. J. 81 (324), 828–846.

Doffkay, R., 2024. Urban Sprawl and Public Transport: A Quantitative Spatial Model for a Railway Policy in Budapest, MSc Transport (Economics Thesis). Technische Universität Dresden, Dresden, Germany.

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

Duranton, G., Puga, D., 2004. Micro-foundations of urban agglomeration economies. In: Handbook of Regional and Urban Economics, vol. 4. Elsevier, pp. 2063–2117.

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

Eliasson, J., Fosgerau, M., 2019. , 'Cost-benefit analysis of transport improvements in the presence of spillovers, matching and an income tax. Econ. Transp. 18, 1–9.

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

Fajgelbaum, P.D., Gaubert, C., Gorton, N., Morales, E., Schaal, E., 2023. Political preferences and the spatial distribuion of intrastructure: evidence from California's high-speed rail, Technical report, NBER Working Paper No. 31438.

Glaister, S., 1974. Generalised consumer surplus and public transport pricing. Econ. J. 84 (336), 849–867.

Graham, D.J., Gibbons, S., 2019. Quantifying wider economic impacts of agglomeration for transport appraisal: existing evidence and future directions. Econ. Transp. 19, 100121.

Gurran, N., Bramley, G., 2017. Urban Planning and the Housing Market, vol. 10. Springer.

Gutiérrez-I-Puigarnau, E., van Ommeren, J.N., 2010. Labour supply and commuting. J. Urban Econ. 68 (1), 82–89.

GVA, 2018. Crossrail property impact and regeneration study, Technical report, Crossrail Ltd. Prepared by GVA for Crossrail Ltd https://learninglegacy.crossrail.co.uk/wp-content/uploads/2018/07/4D-003-crossrail_property_impact_regeneration_study.pdf (Accessed 25 April 2026).

Head, K., Mayer, T., 2014. Gravity equations: workhorse, toolkit, and cookbook. In: Handbook of International Economics, vol. 4. Elsevier, pp. 131–195.

Heblich, S., Redding, S.J., Sturm, D.M., 2020. The making of the modern metropolis: evidence from London. Q. J. Econ. 135 (4), 2059–2133.

Herzog, I., 2024. The city-wide effects of tolling Downtown drivers: evidence from London's congestion charge. J. Urban Econ. 144, 103714.

Hofbauer, J., Sandholm, W.H., 2002. On the global convergence of stochastic fictitious play. Econometrica 70 (6), 2265–2294.

Hörcher, D., Graham, D.J., 2025. Transport in quantitative spatial economics: a review of theory, empirics, and policy simulations. Working Paper, Imperial College, London.

Hörcher, D., De Borger, B., Seifu, W., Graham, D.J., 2020. Public transport provision under agglomeration economies. Reg. Sci. Urban Econ. 81, 103503.

Hunt, J.D., Simmonds, D.C., 1993. Theory and application of an integrated land-use and transport modelling framework. Environ. Plan. B: Plann. Des. 20 (2), 221–244.

Ibraeva, A., de Almeida Correia, G.H., Silva, C., Antunes, A.P., 2020. Transit-oriented development: a review of research achievements and challenges. Transp. Res. a Policy Pract. 132, 110–130.

Jara-Díaz, S., 2007. Transport Economic Theory. Emerald Group Publishing Limited.

Kleven, H.J., Kreiner, C.T., 2006. The marginal cost of public funds: hours of work versus labor force participation. J. Public Econ. 90 (10–11), 1955–1973.

Koster, H.R., 2024. The welfare effects of Greenbelt policy: evidence from England. Econ. J. 134 (657), 363–401.

Koster, H.R., Hayakawa, K., Tabuchi, T., Thisse, J., 2024. High-speed rail and the spatial distribution of economic activity: evidence from Japan's shinkansen. Working Paper, RIETI Discussion Paper 21-E-003..

Kung, K.S., Greco, K., Sobolevsky, S., Ratti, C., 2014. Exploring universal patterns in human home-work commuting from mobile phone data. PLOS ONE 9 (6), e96180.

Laird, J.J., Venables, A.J., 2017. Transport investment and economic performance: a framework for project appraisal. Transp. Policy 56, 1–11.

Lee, K.H., Tan, B.J., 2024. Urban transit infrastructure and inequality. Rev. Econ. Stat. 1–46.

Lowry, I.S., 1964. A Model of Metropolis. Technical report, Rand Corporation, Santa Monica.

Mackie, P., Graham, D., Laird, J., 2011. The direct and wider impacts of transport projects: a review. In: A Handbook of Transport Economics. Edward Elgar Publishing.

Mackie, P., Worsley, T., Eliasson, J., 2014. Transport appraisal revisited. Res. Transp. Econ. 47, 3–18.

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

Miyauchi, Y., Nakajima, K., Redding, S.J., 2025. The economics of spatial mobility: theory and evidence using smartphone data. Q. J. Econ. 140 (4), 2507–2570.
Monte, F., Redding, S.J., Rossi-Hansberg, E., 2018. Commuting, migration, and local employment elasticities. Am. Econ. Rev. 108 (12), 3855–3890.
Parry, I.W., Bento, A., 2001. Revenue recycling and the welfare effects of road pricing. Scand. J. Econ. 103 (4), 645–671.
Pereira, R.H.M., Saraiva, M., Herszenhut, D., Braga, C.K.V., Conway, M.W., 2021. R5R: rapid realistic routing on multimodal transport networks with R^5 in R. Findings. https://doi.org/10.32866/001c.21262
Proost, S., Thisse, J.-F., 2019. What can be learned from spatial economics? J. Econ. Lit. 57 (3), 575–643.
Redding, S.J., Rossi-Hansberg, E., 2017. Quantitative spatial economics. Annu. Rev. Econ. 9, 21–58.
Redding, S.J., Sturm, D.M., 2008. The costs of remoteness: evidence from German division and reunification. Am. Econ. Rev. 98 (5), 1766–1797.
Robson, E.N., Wijayaratna, K.P., Dixit, V.V., 2018. A review of computable general equilibrium models for transport and their applications in appraisal. Transp. Res. a Policy Pract. 116, 31–53.
Santos Silva, J.M.C., Tenreyro, S., 2006. The log of gravity. Rev. Econ. Stat. 88 (4), 641–658.
Severen, C., 2023. Commuting, labor, and housing market effects of mass transportation: welfare and identification. Rev. Econ. Stat. 105 (5), 1073–1091.
Small, K.A., 2012. Valuation of travel time. Econ. Transp. 1 (1–2), 2–14.
Small, K.A., Verhoef, E.T., 2007. The Economics of Urban Transportation. Routledge.
Tikoudis, I., Verhoef, E.T., van Ommeren, J.N., 2015. On revenue recycling and the welfare effects of second-best congestion pricing in a monocentric city. J. Urban Econ. 89, 32–47.
Tsivanidis, N., 2026. Evaluating the impact of urban transit infrastructure: evidence from Bogotá's transmilenio. Am. Econ. Rev. 116 (2), 418–463.
Valentinyi, A., Herrendorf, B., 2008. Measuring factor income shares at the sectoral level. Rev. Econ. Dyn. 11 (4), 820–835.
Venables, A.J., 2007. Evaluating urban transport improvements: cost–benefit analysis in the presence of agglomeration and income taxation. J. Transp. Econ. Policy 41 (2), 173–188.
Vickerman, R., 2017. Beyond cost-benefit analysis: the search for a comprehensive evaluation of transport investment. Res. Transp. Econ. 63, 5–12.
Wangsness, P.B., Rødseth, K.L., Hansen, W., 2017. A review of guidelines for including wider economic impacts in transport appraisal. Transp. Rev. 37 (1), 94–115.
Wardman, M., Chintakayala, V.P.K., de Jong, G., 2016. Values of travel time in Europe: review and meta-analysis. Transp. Res. a Policy Pract. 94, 93–111.
Wegener, M., 2011. Transport in spatial models of economic development. In: de Palma, A. (Ed.), A Handbook on Transport Economics'. Edward Elgar Cheltenham, Glos, pp. 46–66.
Wrigley, E.A., 2011. The Early English Censuses. Oxford University Press.