CPB Netherlands Bureau for
Economic Policy Analysis

Cycling cities: Mode choice, car congestion, and urban structure

This study examines the impact of cycling on urban spatial structure and welfare through the development of a quantitative spatial model that incorporates mode choice and car congestion.
We apply this model to the Netherlands, which is known for its extensive cycling infrastructure.
Eliminating cycling increases commuting times and distances by 14% and 30%, respectively, exacerbates car congestion, and results in a significant reduction in worker welfare.

We also show that removing dedicated cycleways leads to similar, though less pronounced, changes in the spatial economy. Therefore, promoting cycling can help create more compact cities.

CPB Discussion Paper

Tijl Hendrich, Nicole Loumeau, Hans Koster,
Jos van Ommeren

February 2025

https://doi.org/10.34932/qyhd-5t84

Cycling cities:

Mode choice, car congestion, and urban structure*

Tijl Hendrich† Hans R.A. Koster‡ Nicole Loumeau§ Jos van Ommeren¶

December 11, 2024

Abstract — Cycling as a mode of commuting has gained popularity in cities worldwide. This study examines the impact of cycling on urban spatial structure and welfare through the development of a quantitative spatial model that incorporates mode choice and car congestion. We apply this model to the Netherlands, which is known for its extensive cycling infrastructure. Eliminating cycling increases commuting times and distances by 14% and 30%, respectively, exacerbates car congestion, and results in a significant reduction in worker welfare. We also show that removing dedicated cycleways leads to similar, though less pronounced, changes in the spatial economy. Therefore, promoting cycling can help create more compact cities.

Keywords — Cycling, bicycles, car congestion, agglomeration, density, infrastructure, commuting, 15-minute cities

JEL codes — R11, R12, R13, R14.

*Our results are based on our own calculations using non-public microdata from Statistics Netherlands. Under certain conditions, these microdata are accessible for statistical and scientific research. For further information, see: microdata@cbs.nl.

†CPB Netherlands Bureau for Economic Policy Analysis, corresponding author, email: t.j.m.hendrich@cpb.nl; Department of Spatial Economics, Vrije Universiteit Amsterdam.

‡Department of Spatial Economics, Vrije Universiteit Amsterdam, email: h.koster@vu.nl; Tinbergen Institute, CEPR.

§CPB, Netherlands Bureau for Economic Policy Analysis, email: n.loumeau@cpb.nl.

¶Department of Spatial Economics, Vrije Universiteit Amsterdam, email: jos.van.ommeren@vu.nl; Tinbergen Institute.

1 Introduction

In 1900, when biking was a necessity for all but the rich, bicycle ownership in European countries was roughly similar (ANWB, 1900). However, a century later, the Netherlands has emerged as the world's leading cycling nation. With 23 million bicycles in a country of 17.7 million people, bicycle ownership has become almost universal. Approximately 28% of all trips made by bicycle (Ranking Royals, 2023). Denmark follows with 18% of trips, while Sweden and Germany register 12% and 9%, respectively.

In the Netherlands, cycling is also a popular choice for commuting, with approximately 25% of commutes done by bicycle. In its major cities, this percentage is even greater, with most people cycling to their jobs. In contrast, in most cities around the world, commuters hardly ever choose to cycle. One may wonder why. Although countries such as the Netherlands and Denmark benefit from a mild climate and flat terrain, this also applies to numerous other regions, in particular cities located on flat planes (see Figure 1). This includes cities such as Stockholm, Budapest, and London that have natural features that could make them as attractive for cycling as Amsterdam. Although cultural traits may partially explain the high cycling rates in the Netherlands or Denmark, it is clear that past bicycling policies in these countries also played a crucial role in promoting today's bicycle usage.[1]

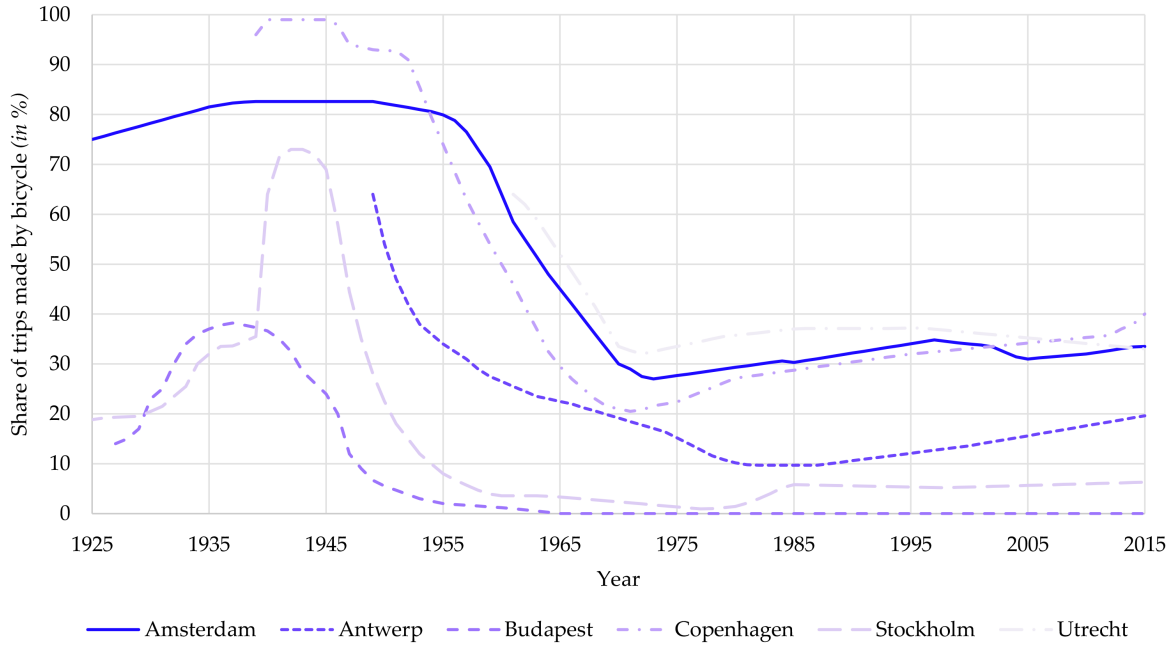

Note: This figure shows the share of trips made by bicycle (in %) over time (years 1925-2015) for six selected European *cycling* cities. The data are from Oldenziel and Albert De La Bruhèze (2011).

FIGURE 1 – BICYCLE USE IN SELECTED EUROPEAN CITIES

[1]The presence of shared bicycle systems in about 300 cities globally suggests that in many other cities mass commuter cycling would be a viable option.

Arguably, a key cycling policy is establishing dedicated lanes that are separated from the main road. Separate cycleways significantly improve cyclist safety and convenience by providing clear and protected space for cycling, free from motor vehicle interference (Thomas and DeRobertis, 2013). In the buildup area of cities, separate cycleways are typically established by reallocating road space previously used by cars, buses and bicycles. As the monetary costs of bicycle lanes are below those of other roads, the primary expense is the opportunity cost of reduced road space for other modes of transport, particularly the car. Separate cycleways are therefore potentially cheap to implement and are believed to address some fundamental challenges facing large metropolitan areas today, including car congestion, urban sprawl, and heavy traffic pollution.

In this paper, we study the effects of bicycle commuting on the spatial structure of cities in the Netherlands. The effects of bicycling as a commuting mode have been largely overlooked in the economics literature. According to urban economics theory, however, one might anticipate that an increase in bicycle commuting in global cities, at a level comparable to the Netherlands, would lead to significant changes in their spatial organisation. This would occur because the marginal commuting costs are highly dependent on the chosen mode of travel and are and therefore a key factor in shaping the spatial structure of cities (Glaeser and Kahn, 2004; Baum-Snow, 2007).

Focusing on the Netherlands offers a distinct advantage as it exhibits significant spatial variation in cycling flows between urban and rural areas. This provides valuable identifying variation for estimating the impact of cycling on urban spatial structure. The high rate of bicycle adoption also enable us to explore compelling counterfactual scenarios, in particular the effects of the separated cycleways. Conducting similar counterfactuals in other countries would require assumptions about increasing bicycle use to levels not yet observed, thereby challenging the external validity of the model.

For a first intuition on the influence of choosing a mode on spatial structure, we turn to the textbook monocentric city model. This model predicts that households who prefer modes with higher marginal costs will reside closer to the city centre (Gin and Sonstelie, 1992; Glaeser, 2008).[2] We illustrate this for three modes in Figure 2: walking, cycling and using the car. In the figure, the use of specific travel modes is influenced by the initial preferences for those modes, as reflected by the intercept of the bid-rent functions. Because car use implies the fixed costs of owning a car and therefore lower bid-rent intercepts, but lower marginal travel costs, the model predicts that at short distances, households prefer to walk or to use the bicycle. Taking the predictions of the monocentric model seriously would mean that removing bicycles and forcing a

[2] Glaeser et al. (2008) highlight that in the US poorer individuals tend to live in city centres because they rely more on public transport, which, due to its higher marginal travel costs, encourages them to relocate to these locations.

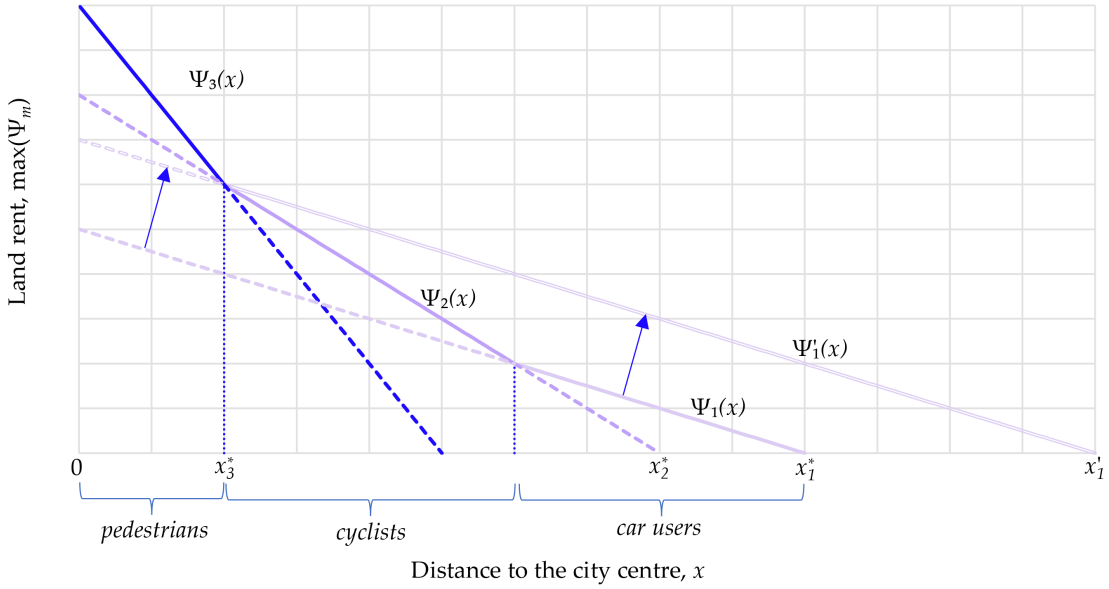

Notes: This figures illustrates the bid-rents for people using different modes are denoted by Ψ_m for each mode $m = 1$ (car), 2 (bicycle), 3 (walk). The land rent consists of the upper envelope of the bid rent functions. Locations are indexed by x, where $x = 0$ is the Central Business District (CBD).

FIGURE 2 – THE MONOCENTRIC CITY MODEL WITH MULTIPLE MODES

switch to cars would lead to greater urban sprawl, as this would shift the bid-rent curve for car users (Ψ_1 to Ψ_1') upward as indicated in the figure. In contrast, if cyclists switch to walking, cities would become more compact. Consequently, the substitution patterns between travel modes is a key determinant of the effect of removing cycling on the spatial economy.

This also raises the question how car congestion interferes with this conclusion. Does car congestion amplify or reduce the effect of removing cycling on the spatial Structure? Maybe paradoxically, it appears that car congestion reduces the effect of removing cycling on the spatial structure. The figure above clearly demonstrates this point. If car congestion increases, it leads to a rise in the marginal cost of car travel. In the particular case where the marginal costs of car travel equal those of bicycle travel, there would be no effect on the spatial structure.

Despite its elegance, this monocentric framework may not capture reality well, as both residential and job locations are not fixed but endogenous. Additionally, the substitutability between modes depends heavily on workers' willingness to commute longer. General equilibrium effects, including changes in housing supply and agglomeration economies, further complicate the story. Consequently, the extent to which mode choice – particularly cycling – influences urban spatial structure remains unclear.

Therefore, we develop a spatial quantitative model based on Ahlfeldt et al. (2015) and Heblich et al. (2020) among others to study effects of cycling on spatial structure . Our model includes commuting, choosing the location of residence and work, car congestion and agglomeration economies. Similarly to Tsivanidis (2023), we explicitly incorporate choices of travel modes that

are consistent with household utility maximisation.

In our model, we allow for four modes: bicycle, walking, public transport and car. We allow household preferences for travel mode and residential-workplace location to be correlated by introducing a nested choice structure. This structure can be interpreted as if workers first select their residential-workplace location, subsequently decide between using a car or other travel modes, and then choose whether to walk, cycle or use public transport. Car use is affected by congestion within the area of the workplace, which we model as a 'bathtub filled with traffic' (Arnott, 2013; Fosgerau, 2015). Car drivers then adjust their route choices to avoid congestion in the area. This adjustment process results in an equilibrium where congestion is evenly distributed within the workplace area and a stable relationship between car density and speed across the area emerges. Cyclists value the separate cycling infrastructure as it decreases their travel costs. Gains in safety are not explicitly modeled.

Next, we estimate the model for the Netherlands. Estimation of spatial quantitative models, rather than calibration, offers several advantages, as model parameters are derived directly from relevant data rather than being borrowed from unrelated contexts. In addition, some parameters are difficult to calibrate due to the lack of international benchmarks. For instance, there are no proper benchmarks for spatial decay in commuter flows for various travel modes. Furthermore, some relationships are likely country-specific. For example, the relationship between speed and traffic density has not been estimated for Dutch cities, which is problematic as it strongly depends on the quality of the underlying infrastructure (Akbar et al., 2023). To illustrate, in the Netherlands, about half of all car kilometres are driven on highways, much higher than in most other countries. We rely on unique microdata on commuting mode choices as well as wages, residential and workplace location choices from *Statistics Netherlands*.

We then utilise our model to perform three counterfactual experiments. In the first experiment, we exclude bicycles from the set of choices (and re-assign separate cycleways to cars in dense urban areas). We find that the removal of bicycles results in a substantial 20 percentage point increase in car commuting. This increase in car use leads to longer average commuting distances, up by 30%, and a 2% decrease in car speeds due to increased traffic congestion. Hence, lower per-km costs of driving contribute to increased urban sprawl. Eliminating bicycle commuting reduces worker welfare by about 6%. These results suggest that cycling not only helps to maintain more compact cities, it may even be the case that bicycle-friendly policies increase welfare.

The second experiment is arguably more realistic, and as a policy more explicit, as it removes separate cycleways, forcing many cyclists to take detours, as not all roads are suitable for cyclists.

This experiment mimics the situation in the US, where cycling often requires avoiding large roads and taking detours using residential streets. We show that the share of bicycle commuters decreases by about 20%, and welfare decreases by 1.2%. As before, making cycling less attractive leads to longer commute times and distances, with increases of 1.2% and 7%, respectively.

The final experiment assumes that bicycle speeds will increase in the future, reflecting the growing adoption of electric bicycles (Yin et al., 2024). Our results show that a 10% increase in bicycle speeds leads to an 11 percent rise in cycling mode share, drawing users away from cars, walking and public transport. As a result, the welfare of workers improves by 0.8%, and the residential densities in inner cities increase by approximately 2%.

Related literature and contribution. With this paper, we contribute to three strands of literature, namely the literature on transport mode choice, the rapidly growing body of literature on Quantitative Spatial Equilibrium (QSE) models, and an emerging literature on so-called 15-minute cities. Meanwhile, we increase the depth and level of detail of the analysis by using a rich universe of Dutch microdata along with the complete Dutch transport networks including roads, cycleways, and public transport.

Our work first contributes to a broad body of literature on mode choice, specifically focusing on the decision to cycle.[3] Rietveld and Daniel (2004) demonstrate that variations in bicycle use in Dutch municipalities are significantly influenced by physical factors such as altitude and city size, along with demographics. The relative speed and parking advantages of bicycles over cars also explain spatial variation in cycling to a large extent. Buehler and Dill (2016) confirm a positive correlation between the quality of bicycle infrastructure and cycling levels. Heinen et al. (2010) identify additional factors affecting bicycle use, including infrastructure quality, land use, and weather conditions, noting that car ownership reduces cycling while bicycle ownership promotes it.

Our contribution to the mode choice literature is two-fold. First, we model mode choice using a nested logit approach based on revealed preference data, relaxing the IIA assumption common in many studies. Furthermore, stated choice experiments are often plagued by hypothetical bias, making the use of revealed preference data clearly superior (Loomis, 2011). Second, we integrate the choice of cycling with residential and workplace location decisions, acknowledging that a preference for cycling requires living relatively close to work. This approach allows for endogenous mode-location choices, which contrasts to the transport literature, which assumes

[3] See Nelson and Allen (1997); Aultman-Hall et al. (1998); Handy et al. (2002); Cervero and Duncan (2003); Pucher and Dijkstra (2003); Saelens et al. (2003); Rietveld and Daniel (2004); Krizek and Johnson (2006); Pucher and Buehler (2006); Titze et al. (2008); Akar and Clifton (2009); Cervero et al. (2009); Dill (2009); Pucher et al. (2010); Xing et al. (2010); Winters et al. (2011); Beenackers et al. (2012); Broach et al. (2012); Buehler (2012); Schoner and Levinson (2014); Ton et al. (2019); Rayaprolu et al. (2020); Ton et al. (2020).

that residential and workplace locations are fixed (Small et al., 2007).

In the realm of the QSE literature, our study aims to advance the field in three significant ways. Our paper builds on Ahlfeldt et al. (2015), who develop a model that includes commuting flows, choice of workplace and residential location and agglomeration economies (see for an overview Redding, 2023).[4] Recent spatial quantitative models have also considered a more explicit modelling of shopping trips, firm location choice, and include land use regulation, among others (see, respectively, Miyauchi et al., 2021; Monte et al., 2018; Dericks and Koster, 2021).

As a first contribution, we explicitly analyze the equilibrium effects of infrastructure changes on mode choice, while we examine the consequences of common assumptions such as Independence of Irrelevant Alternatives (IIA) and mode-specific commuting time cost (Kouwenhoven et al., 2014) for the counterfactual outcomes. In doing so, we map the triad of interactions between mode choice, location choice and car congestion in a spatial equilibrium framework.[5] To date, only a limited number of studies applying spatial quantitative models have incorporated mode choice (Tsivanidis, 2023; Severen, 2023; Koster, 2024; Koster et al., 2024). The study most closely related to ours is Tsivanidis (2023) on Bogotá, Colombia. It includes two modes: car and public transport, while relaxing the IIA assumption, with a particular focus on the TransMilenio, a Bus Rapid Transit service. In this study, commuters choose the mode that minimises their commuting costs, given their chosen residential and workplace locations.[6] The work of Tsivanidis (2023) is a significant contribution to the literature; however, it relies on limited aggregate information on mode choices from a small survey and therefore does not have information on location pair-specific mode choices. As a result, it is silent on the (spatial) shift in mode choice at the location pair that occurs because of changes in mode-specific travel costs. Similar limitations apply to the study by Severen (2023), which investigates the effects of LA Metro Rail on commuting flows and housing and labour market outcomes, but not on choice of modes. The study by Koster (2024) focuses on commuters in the UK and includes mode choice among three options (car, public transport, other) and accounts for traffic congestion,

[4]To the best of our knowledge, the only spatial equilibrium model that has been estimated for the Netherlands is Teulings et al. (2018), who show that investments in public transport infrastructure have heterogeneous effects on location choices, significantly boosting the welfare of the highly educated. They model the location choice of homes, workplaces, and modes as sequential. The paper clearly illustrates the outstanding quality of the Dutch micro-data, but their model does not allow for an inversion of the amenities and productivities so that the observed equilibrium does not match the predicted equilibrium. Moreover, Teulings et al. (2018) do not take into account the choice of firm location and assume a uniform dwelling size. We therefore take the more standard approach as outlined in Ahlfeldt et al. (2015).

[5]Although the emerging QSE literature shares strong connections with international trade studies (*e.g.*, Eaton and Kortum, 2002)), it is less rooted in research on mode choice and travel behaviour. Our intent is to apply insights from the transport economics literature to the QSE literature.

[6]In the study, car ownership is also explicitly modeled, which introduces complications, as it is highly endogenous. In high-income countries such as the Netherlands, introducing this complication has few advantages, as income levels are high enough for all workers to buy a car.

but does not relax the IIA assumption. This implies that mode choice substitution is imposed by the structure of the model and is not estimated using data. Koster et al. (2024) examine the commuting choices between car and train for Japanese workers, focusing on only two travel modes, which results in a simplified nested structure.[7]

Second, to our knowledge, none of the studies in the QSE literature consider bicycles as a mode of transport or examine the substitution between different modes, which has consequences for commuting flows and the urban spatial structure. Our study fills this gap by incorporating bicycle use, which is (potentially) a close substitute for walking, public transport, and car use for shorter distances (that is, *within* cities), but not for longer distances. Consequently, bicycles offer an interesting alternative for a specific share of commuting trips, but certainly not all, making the trade-offs in mode choice and location choice more nuanced. By estimating a marginal cost of travel time per mode, we aim to provide a deeper understanding of commuting behaviour and mode choice in urban areas in the presence of equilibrium effects. We will demonstrate that using mode-specific time costs is important, as failing to do so significantly biases the aggregate mode shares and therefore underestimates welfare implications of cycling infrastructure and cycling in general.

Third, we explicitly analyse the interaction between car congestion, transport infrastructure provision and mode choice. A significant recent contribution by Allen and Arkolakis (2022) accounts for the longer commute times on busier routes, demonstrating the importance of considering traffic congestion in predicting the returns on investments in transportation infrastructure. We do not adopt their approach because we lack link-specific car traffic data on Dutch inner city roads. Instead, we follow the approach introduced by Koster (2024), who models congestion parsimoniously as a 'bathtub' (see Arnott, 2013; Fosgerau, 2015). In this setting, the speed between residential and workplace locations depends on the density of workers who commute by car.

Finally, with the current research we aim at contributing to the emerging field of research on compact cities and so-called 15-minute cities. 15-minute cities refer to an urban planning concept in which residents can access most daily necessities, including work, within a 15-minute ride from their homes (Allam et al., 2022). Bicycles play a crucial role by providing an environmentally-friendly mode of transport that reduces the dependency on cars, reduces emissions, and promotes health (Moreno et al., 2021). However, there is yet no evidence to what extent bicycles actually impact commuting trips and reduce commute distances and times. This is where our paper makes a contribution, shedding light on the role of cycling in urban mobility.

[7]Additionally, due to a lack of detailed data on mode choice, they derive mode shares solely based on the structure of their model rather than on empirical observations.

The paper is structured as follows. Section 2 describes the context, outlines the various data sources, and presents some basic descriptive statistics. In Section 3 we outline our model. Section 4 turns to the estimation, followed by the results of several counterfactual experiments in Section 5. Section 6 concludes.

2 Research context and data

2.1 Cycling in the Netherlands

One of the key determinants of cycling use is the presence of a high-quality cycling infrastructure (Buehler and Dill, 2016). The first bicycle path was opened between the Dutch villages of Baarn and Laren in 1914 and was an instant success (ANWB, 2024). However, the most significant increase in cycling path construction occurred in the 1970s and 1980s, with approximately 25,000 km of cycling paths built during that period. The cyclist death toll doubled between the 1950s and 1970s, prompting widespread protests by the pressure group *Stop de Kindermoord* ("*Stop child murder*"). The protesters occupied dangerous intersections, organised bicycle demonstrations, and closed streets for children to play (Oldenziel and Albert De La Bruhèze, 2011; Bruno et al., 2021). In response, the Dutch government initiated a nationwide plan to expand and improve the bicycle path network, enhancing safety and comfort for cyclists.

In addition to developing an extensive cycling infrastructure within and between cities, the government implemented policies to improve bicycle friendliness. These included introducing traffic rules that prioritise cyclists, providing financial support and tax breaks to the bicycle industry to foster production and innovation, and promoting the construction of widespread bicycle parking facilities. The government also introduced laws to promote cyclist safety. In particular, if there is a collision between a car driver and a cyclist, the former is fully accountable for damages when they were at fault. However, when the cyclist was at fault, the driver of the car still has to pay half of the damage. In addition, municipalities can be held accountable if they neglect road maintenance and upkeep or allow unsafe infrastructure to exist.

Today, cycling infrastructure is ubiquitous in the Netherlands. On residential streets within neighbourhoods with little traffic, car and bicycle traffic share the road. On-road cycling lanes separate cyclists from the road traffic with a dashed or straight line (see the left and middle panel of Figure 3). Busier roads are constructed with separate cycleways (right panel of Figure 3). This is a salient difference to developments in other major European countries, where cycling infrastructure is lacking or the emphasis is on on-road cycling lanes. In the Netherlands, the separate cycleways played an important role in physically separating traffic streams and thereby

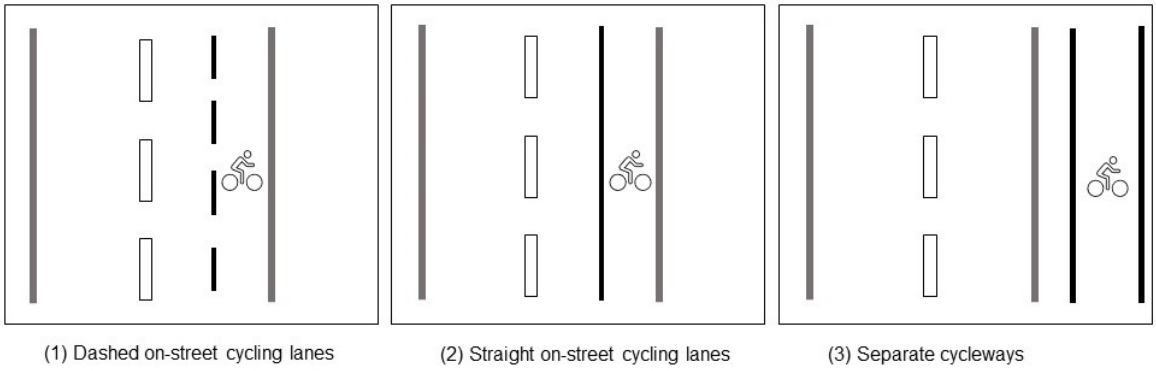

Notes: This figure provides a graphical illustration of the three types of cycling infrastructure in the Netherlands. On-street cycling lanes are either marked with a dashed line (left panel, (1)), indicating that cars and bicycles share the lane, or with a straight line (middle panel, (2)), indicating that cars are not supposed to cross that lane. Separated cycleways (right panel, (3)) are typically physically separated from motor traffic and for cycling exclusively.

FIGURE 3 – THREE TYPES OF CYCLING INFRASTRUCTURE

creating a safe and comfortable cycling environment. The Netherlands boasts 30,997 km of dedicated cycleways along 43,967 km of larger roads for through traffic (ECF, 2023). This means that most major roads (70%) are accompanied by separate cycleways.

To ensure uniformity of cycling infrastructure layout and safety measures, countrywide design guidelines and principles for cycling infrastructure have been formulated and updated by a consortium of stakeholder representatives since the 1970s (CROW, 2017).

2.2 Commuting patterns

Based on the 2016-2022 annual *National Travel Survey* (`OViN`), provided by Statistics Netherlands, we analyse the commuting patterns of a random sample of the Dutch population. The survey provides information on travel mode, as well as residence and work location at the neighbourhood level for in total 48,323 observations. A large proportion of workers commute by car (65%), while 25% use bicycles. For shorter distances (<5 km), most opt for cycling (55%), followed by car use (29%). The average commuting time stands at 26 minutes in 2016.

Public transport, which consists mainly of train travel, is used mostly for long-distance commutes *between* cities in the Netherlands.[8] A major difference compared to other countries is that the Netherlands has many cities located relatively close to each other, which makes intercity commuting attractive, especially if cycling speeds increase in the future. For example, the distance between the second and third largest cities in the Netherlands - Rotterdam and The Hague - is only 20km.

[8]Approximately two-thirds of public transport commuters use the train. Of these, a significant portion – around 32% – cycle to the train station, while 8% use bicycles to travel from the train station to work (Martens, 2004). In our counterfactual analysis, we ignore multimodal trips, likely underestimating the role of cycling in shaping commuting patterns and influencing the urban spatial structure.

TABLE 1 – SUMMARY STATISTICS COMMUTING

	Car	Bicycle	Walking	Public transport	Total
Panel A: All trips					
Mode share	0.65	0.25	0.04	0.06	1
Commuting time (in min)	27	18	10	53	26
Distance (in km)	24	4	1	31	18
Median speed (in km/h)	45	14	5	31	
Panel B: Within city trips, < 5km					
Mode share	0.29	0.55	0.14	0.02	1
Commuting time (in min)	9	12	10	21	11
Distance (in km)	3	2	0.8	3	2
Median speed (in km/h)	18	12	5	9.4	12
Share of total trips	0.14	0.68	0.99	0.08	0.31

Notes: The data are sourced from the `OViN` for the years 2016-2022 survey with a total number of observations of 48,323. We only consider work commuting trips, *i.e.*, recreational trips or shopping trips are excluded. Commuting time and commuting distance are self-reported. Panel A reports summary statistics for all trips, while Panel B reports the same summary statistics for trips shorter than 5km. Short trips start in the majority of cases (>90%) in highly urban neighbourhoods (*i.e.*, more than 1000 addresses per km^2).

We observe that in 2017, the total number of cross-border commuters to Germany and Belgium was only 22,340 workers, representing approximately 0.002% of the entire working population. Therefore, border effects are not a significant concern and we ignore cross-border commutes in what follows.

2.3 Data description

Neighbourhoods, population and workplaces. Our main data source is the universe of Dutch 2016 microdata compiled by Statistics Netherlands. All data are based on administrative sources, unless otherwise indicated. Our unit of analysis is the neighbourhood, defined by Statistics Netherlands. We have information on 12,782 neighbourhoods, which on average contain approximately 600 households and span the country.[9] To ensure that short-distance trips, including most bicycle commutes, are captured correctly, we use small spatial units for analysis.[10] Neighbourhoods are designed to have a similar population size, meaning that in rural areas they cover larger geographical spaces, while in densely populated cities they span smaller areas.

We use the Dutch civil register (`GBA`) and address register (`GBAADRESOBJECTBUS`) to obtain the place of residence of each worker at the neighbourhood level. On the workplace side, we extract

[9] We drop a small number of neighbourhoods on the Wadden Islands, since we lack information about travel time by ferry.

[10] Although we have a very granular spatial setting, our calibration procedure discussed later on is based on the observed commuting time and travel mode, rather than on observed commuting shares, which avoids the issue that idiosyncratic choices affect equilibrium outcomes that result in very poor counterfactual predictions, as discussed in Dingel and Tintelnot (2020).

firm branch locations from the national firm register (ABR). If a firm has multiple branches, we set the workplace to the location closest to the worker's residence following Gaigné et al. (2022).

Mode choice and travel times. We use the years 2016-2022 OViN survey for information on mode choice per neighbourhood combination. Within the OViN survey we consider only work commuting trips, *i.e.*, recreational trips or shopping trips are excluded. For the estimation of key parameters in the model, we impute mode choices for all neighbourhood combinations using the OViN sample and applying machine learning methods[11]

We calculate travel times between neighbourhood pairs for four travel modes.[12] Travel times are determined along the road, cycling, and pedestrian networks using the complete Dutch road infrastructure from the *Open Street Map* (OSM) database combined with assumptions on average travel speeds per mode and network layers. The OSM database differentiates between different types of roads (*e.g.*, motorways, primary, secondary, and tertiary roads) and the type of cycling infrastructure (on-road cycling lanes and separate cycleways). We use average free-flow travel speeds per mode and type of road from the dodgr graph routing implementation (Padgham, 2019).[13] We determine the congested car speeds using data sampled from the Google Maps API based on historic traffic levels (see Appendix A.4).

For public transport, we use the 2018 operating schedule and timetables of public transport providers in *General Transit Feed Specification* (GTFS) format from Open Geo. The GTFS data contain all public transport stops, arrival and departure times of buses, trams, metros, trains, and ferries. The routing algorithm which computes the fastest route allows for transfers within and between modes (*e.g.*, from bus to tram).

Wage data. We obtain yearly wages for all Dutch workers from income registry data SPOLIS (*Salaris Polis Administratie*). We discard workers who receive unemployment and disability benefits, state pensions, firm owners, and temporary workers since they either do not have a fixed work location or their travel behaviour is likely to be governed by factors outside of our model. We also drop workers who work less than 8 hours per week, as they are unlikely to commute to work multiple times a week on a regular basis.

Floor space prices. We estimate floor prices conditional on housing characteristics using housing transaction data by the Dutch Association of Real Estate Brokers (NVM), as in Teulings et al.

[11] These imputed mode shares are used for estimating the congestion elasticity as well as agglomeration economies in productivity and amenity, see Sections 4.3 and 4.4. More details on the imputation method are provided in Appendix A.5.

[12] In our mode choice and gravity estimations, we exclude trips longer than 120 minutes, since these are likely to be affected by factors outside of our model.

[13] We present details on the travel speeds by network type and the routing procedure in Appendix A.

(2018) and Gaigné et al. (2022). The NVM data, covering the large majority of owner-occupied housing transactions, include transaction prices, lot sizes, interior floor space size (m^2), location and numerous housing attributes (house type, number of rooms, construction year, garden, state of maintenance, central heating, and listed building status).

Historic data. We use historic population data from 1900 as an instrumental variable to estimate some parameters of the model. This population data comes from *NLGIS* for municipalities in 1900, which were smaller and comparable to the size of a large neighbourhood today. The local population distribution is imputed by mapping buildings and assuming a uniform population per building within each municipality based on detailed data on land use provided by Knol et al. (2004).[14]

3 The model

This section introduces the theoretical framework that is amenable to analyzing urban structure. We adapt the model of Ahlfeldt et al. (2015) in the spirit of Tsivanidis (2023) by introducing a nested mode choice structure. Furthermore, we allow for congestion that endogenizes car commuting time between locations.

3.1 The economy

The economy consists of a finite number of pairwise-connected locations $i = 1, \ldots, I$. Each location has an area of developed land L_i and floor space F_i, which can be used for commercial (subscript M) and residential (subscript R) purposes. The share of developable land devoted to transport infrastructure is denoted by ι_i, hence $(1 - \iota_i)L_i$ is available for residential and commercial development. The economy is populated by H workers, who are perfectly mobile in the housing and labour market and commute between locations by means of transport mode m. Workers either commute by car, bicycle, public transport or walking.[15]

Commuting costs are modeled as iceberg costs and depend on the generalised commuting time, which again depends on the utility offered by available mode choices, which depend on the mode-specific commuting time. Traveling by car is subject to congestion, which increases with the number of car commuters and decreases with the number of car lanes at the workplace location. We feature the trade-off in the allocation of urban land to car or bicycle infrastructure

[14] For the 1900 land use maps, Knol et al. (2004) scanned and digitised historic maps into 50×50 metre grids, classifying them into 10 categories, such as built areas, water, sand, and forest. We aggregate these categories into three broad groups: built-up areas, open space, and water bodies.

[15] Public transport includes commuting by bus, tram, metro and train. In practice, some commuters combine different modes, but for the sake of simplicity, we focus on the main mode of transport only.

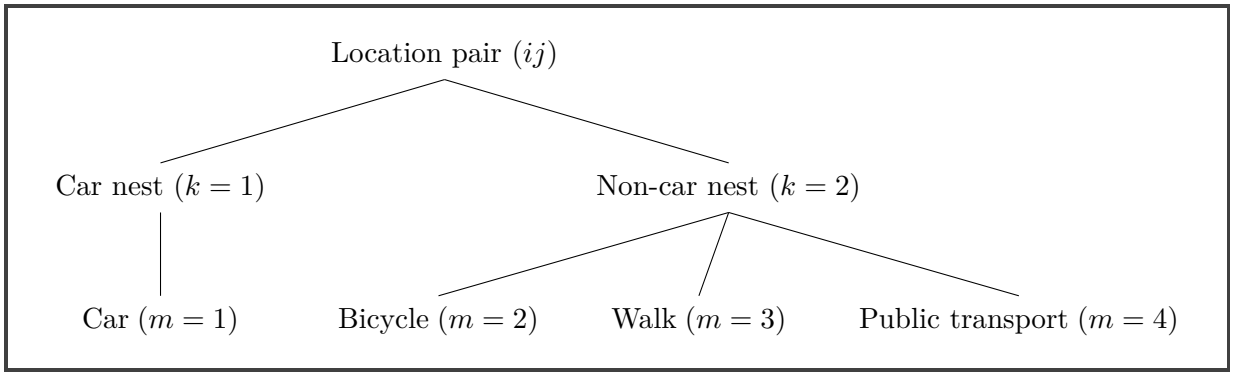

Notes: This figure is a graphical illustration of the nested choice structure in our model. Workers first select their residential-workplace location, subsequently decide between using a car or other travel modes, and then choose whether to walk, cycle, or use public transport.

FIGURE 4 – NESTED CHOICE STRUCTURE

by assuming that the share of land devoted to transport infrastructure, ι_i is fixed. Hence, for each added bicycle lane, the car lane capacity reduces accordingly. Increasing bicycle infrastructure, therefore, affects car congestion through two channels: a reduction in road capacity leading to more car congestion and mode-substitution away from cars reducing congestion. The total effect on car congestion depends therefore on the relative strength of either force.

3.2 Worker utility: amenities, location and travel mode

Each worker o derives utility from consumption c_{ijo} of a single composite good, from living in a dwelling of size l_{ijo}, and from enjoying residential amenities B_i and workplace amenities C_j.[16] After observing their idiosyncratic location preference shock, ξ_{ijo}, workers choose a location pair that maximizes their utility which takes on the following Cobb-Douglas form:

$$U_{ijo} = B_i C_j \left(\frac{c_{ijo}}{\beta}\right)^{\beta} \left(\frac{l_{ijo}}{1-\beta}\right)^{1-\beta} \xi_{ijo}, \tag{1}$$

where β governs the share of non-housing expenditure in the utility. The idiosyncratic component, ξ_{ijo}, is drawn from a Fréchet distribution with shape parameter $\varepsilon > 1$ and scale parameters $B_i, C_j > 0$.[17] Workers earn a wage w_j at their workplace j and pay a rent $p_i l_{ijo}$ in i, where p_i is the floor price per unit in i. Commuting to work is costly and proportional to the wage. Consequently, workers earn an income net of commuting costs of $\bar{w}_j = w_j/d_{ij}$, where $d_{ij} \geq 1$ reflects commute costs.

Workers select their residence-workplace-mode combination ijm in a sequential manner as illustrated in Figure 4. More specifically, each worker o chooses a location pair ij that

[16] Workplace amenities are exogenous to the employer and the worker. Important examples include the presence of shops, parks and restaurants, which can be visited during lunch or after work.

[17] Formally, the probability that individual o in with location-choice ij draws an idiosyncratic preference smaller than z is given by $\Pr(\xi_{ijo} \leq z) = \exp(-B_i C_j \xi_{ijo}^{-\varepsilon})$, where $\varepsilon > 1$.

maximizes utility. Given the location choice ij, the worker chooses the mode of transport m in two steps. First, the worker chooses between two mode nests k. Here, $k = 1$ refers to choosing the car nest.[18] $k = 2$ refers to the non-car nest. The latter nest contains the choice between bicycle use, walking and public transport. So, commuters have a Nested Logit travel demand (McFadden, 1978), structured in a car nest, $\mathbb{M}_1 \equiv \{\text{Car}\}$ and a non-car nest, $\mathbb{M}_2 \equiv \{\text{Walk, Bicycle, Public transport}\}$.

3.3 Location choice

A worker o chooses a residence i and a workplace j that maximizes the indirect utility. From (1), the indirect utility, u_{ijo} can be derived as:

$$u_{ijo} = \frac{B_i w_j}{d_{ij} p_i^{1-\beta}} \xi_{ijo}. \tag{2}$$

The indirect utility is proportional to the idiosyncratic Fréchet-distributed preference shock. As a consequence, it is also Fréchet-distributed and we can derive the probability of choosing a certain ij combination, which takes the well-know gravity form:

$$\pi_{ij} = \frac{a_i \omega_j d_{ij}^{-\varepsilon}}{\sum_{r=1}^{I} \sum_{s=1}^{I} a_r \omega_s d_{rs}^{-\varepsilon}}, \tag{3}$$

where transformed wages ω_j and transformed amenities a_i are defined as

$$a_i = B_i p_i^{\varepsilon(\beta-1)} \quad \text{and} \quad \omega_j = C_j w_j^{\varepsilon}. \tag{4}$$

We define the commuting costs between ij as $d_{ij} = \exp(\eta t_{ij}) \geq 1$, where $t_{ij} \geq 0$ is the generalised commuting time between i and j that accounts for the mode choice set within ij and $\eta > 0$ is a cost parameter that translates generalised commuting time into costs.

We will demonstrate in Section 3.4 that the above-assumed nested structure of mode choice implies that the generalised commuting time between two locations depends on the nest-specific generalised commuting times between these locations, which we will refer to as the *nest-specific commuting time*. More specifically, given two nests $k = 1, 2$, the generalised commuting time t_{ij} is a standard log-sum:[19]

[18] Typically, commuters who travel by car also own a car. Accordingly, car ownership is a consequence of choosing to commute by car.

[19] Generalised commuting time reflects the disutility of commuting between 2 locations, and therefore does not have an intuitive order of magnitude. Therefore, we scale t_{ij} with $\frac{\psi_1}{\kappa_1}$, similar to Tsivanidis (2023).

$$t_{ij} = \log\left[\sum_{k=1}^{2} \exp\left(\psi_k t_{k|ij}\right)\right], \tag{5}$$

where $t_{k|ij}$ is the nest-specific commuting time, which we define further in the next subsection. Conditional on residing in i, the probability of working in j is:

$$\pi_{ij|i} = \frac{\omega_j d_{ij}^{-\varepsilon}}{\sum_{r=1}^{I}\sum_{s=1}^{I} \omega_s d_{rs}^{-\varepsilon}}. \tag{6}$$

Similarly, the probability of residing in i, conditional on working in j is:

$$\pi_{ij|j} = \frac{a_i d_{ij}^{-\varepsilon}}{\sum_{r=1}^{I}\sum_{s=1}^{I} a_r d_{rs}^{-\varepsilon}}. \tag{7}$$

The reservation utility, which equals workers' welfare, equals:

$$\bar{u} = \mathbb{E}[u] = \Gamma\left(\frac{\varepsilon - 1}{\varepsilon}\right)\left[\sum_{i=1}^{I}\sum_{j=1}^{I} B_i C_j \ \left(d_{ij} p_i^{1-\beta}\right)^{-\varepsilon} \ (w_j)^{\varepsilon}\right]^{\frac{1}{\varepsilon}}, \tag{8}$$

where $\Gamma(\cdot)$ is the gamma function.

3.4 Mode choice

Once a worker has chosen their residence-workplace pair ij, they observe their idiosyncratic preferences across modes ϕ_{ijmo}. Workers solve the commuter's mode choice problem by choosing the mode with the lowest generalised commuting time given ij:

$$\min_{m \in M}\{t_{ijm}\}. \tag{9}$$

We define the mode-specific generalised commuting time as a linear function of commuting time τ_{ijm} (which we observe):

$$t_{ijmo} = b_m + \kappa_m \tau_{ijm} + \phi_{ijmo}, \tag{10}$$

where b_m is a preference shifter for mode m and κ_m denotes the mode-specific sensitivity to commuting time. In the following, we refer to κ_m as the mode-specific commuting costs. For Nested Logit travel demand, the idiosyncratic preferences across modes ϕ_{ijmo} are drawn from a Generalised Extreme Value distribution (GEV):

$$F(\phi_{ijmo}) = \exp\left(-\sum_{k=1}^{2}\left(\sum_{m\in\mathbb{M}_k}\exp\left(-\frac{\phi_{ijmo}}{\psi_k}\right)\right)^{\psi_k}\right), \tag{11}$$

where ψ governs the within-nest correlation.[20] Standard results for GEV distributions allow us to write the probability of choosing a mode m conditional on a location-pair ij as two conditional probabilities:

$$\pi_{m|ij} = \pi_{k|ij} \times \pi_{m|kij}. \tag{12}$$

We can express the probability of choosing nest k given ij as:

$$\pi_{k|ij} = \frac{\exp\left(\psi_k t_{k|ij}\right)}{\sum_{n=1}^{2}\exp\left(\psi_n t_{n|ij}\right)} \tag{13}$$

In (13), $t_{k|ij}$ is the nest-specific commuting time of nest k within the location pair ij. Since the car nest is degenerate and contains only one alternative, we can write:

$$t_{k|ij} = \begin{cases} \left(\frac{b_1}{\psi_1} + \frac{\kappa_1}{\psi_1}\tau_{ij1}\right), & \text{if } k = 1 \\ \log\sum_{m\in\mathbb{M}_2}\exp\left(\frac{b_m}{\psi_2} + \frac{\kappa_m}{\psi_2}\tau_{ijm}\right), & \text{if } k = 2 \end{cases} \tag{14}$$

Consequently, the probability of choosing mode m given a nest k can be expressed as follows:

$$\pi_{m|k,ij} = \begin{cases} 1 & \text{if } k = 1 \\ \dfrac{\exp\left(\frac{b_m}{\psi_2} + \frac{\kappa_m}{\psi_2}\tau_{ijm}\right)}{\sum_{n\in\mathbb{M}_2}\exp\left(\frac{b_n}{\psi_2} + \frac{\kappa_n}{\psi_2}\tau_{ijn}\right)}, & \text{if } k = 2 \end{cases}. \tag{15}$$

Knowing the probability of choosing a location pair ij in equation (3), the probability of choosing a nest k in equation (13) and a mode m (15), we can derive the unconditional probability of choosing any given ijm combination as $\pi_{ijm} = \pi_{ij} \times \pi_{k|ij} \times \pi_{m|k,ij}$.

3.5 Car congestion

We allow car commuting times along the road network to be endogenous due to congestion. We model car congestion as a 'bathtub' that fills with traffic (Arnott, 2013; Fosgerau, 2015). Drivers continuously adjust their route choices to avoid heavily congested roads, leading to an equilibrium where congestion is evenly distributed throughout the area. This results in a stable

[20] Note that $1 - \psi_k^2$ is the correlation between the idiosyncratic shocks, ϕ_{ijmo}, for any two modes within the same nest.

relationship between the car density and average speed across the area considered (Daganzo, 2007; Daganzo et al., 2011). Hence, congested commuting time by car ($m = 1$), τ_{ij1}, is assumed to be proportional to free-flow travel time, τ^f_{ij1}, and a function of the density of car traffic at the workplace, Λ_{Mj}, as follows:

$$\tau_{ij1} = \tau^f_{ij1} T_{Ri} T_{Mj} \exp(\lambda_1 \Lambda_{Mj} + \lambda_2 (\Lambda_{Mj})^2), \tag{16}$$

where T_{Ri} and T_{Mj} are location-specific congestion characteristics at the residence and the workplace, respectively, and where λ_1 and λ_2 determine the marginal effect. The semi-elasticity of car density to commuting time is then equal to $\lambda_1 + 2\lambda_2 \Lambda_{Mj}$. Note that if $\lambda_2 = 0$, then we have a standard exponential relationship, which implies that the relationship between commuting time and density is convex. Our assumption that car congestion happens predominantly at the workplace is consistent with the notion that workplace locations are more spatially concentrated than residence locations.

The total number of car commuters in j equals the total number of commuters to the workplace j times the probability of commuting by car, $H_{Mj1} = \sum_i \pi_{k=1|ij} \times H_{Mij}$. We aim to take into account that workers are not only affected by congestion at the workplace but also at nearby locations. Consequently, the density of car traffic at workplace j, which captured the ratio of car commuters to road lanes, is spatially weighted and written as:

$$\Lambda_{Mj} = \sum_{i=1}^{I} \frac{H_{Mi1}}{R_{i1}} \exp(-\kappa_1 \tau_{ij1}), \tag{17}$$

where κ_1 denotes the car commuting cost parameter, which governs the spatial decay of car density and R_{i1} is the number of road lane kilometers in location i.

3.6 Infrastructure land, car lanes and bicycle lanes

The total amount of land used for infrastructure in each location is assumed to be fixed and given by $\iota_i L_i$. For simplicity we assume that $\iota_i L_i$ only contains bicycle lanes and car road infrastructure so we ignore land used for public transport infrastructure and sidewalks. Let R_i be the total amount of lanes, then:

$$\iota_i L_i = R_i = \ell_1 R_{i1} + \ell_2 R_{i2}, \tag{18}$$

where R_{i1} denote the km of car ($m = 1$) lanes and R_{i2} capture the km of bicycle ($m = 2$) lanes. ℓ_1 and ℓ_2 are aggregate scale parameters that transform car and bicycle lanes into land used for infrastructure. Since $\iota_i L_i$ is assumed to be fixed, adding a bicycle lane reduces the space available for existing car lanes. This assumption is realistic, especially in inner cities, where

expanding roadways is typically challenging due to space constraints.

Investments in cycling infrastructure that increase the number of bicycle lanes affect car congestion in two ways. First, additional cycling lanes decrease the number of road lanes. Second, as cycling becomes more attractive, workers will substitute away from car commuting, which decreases the number of car commuters. The overall effect of bicycle infrastructure investments on car congestion therefore depends on the relative strength of either impact.

3.7 Production and agglomeration

Firms produce a single final good, which is costlessly traded between locations, and is chosen as the numéraire. Production occurs under perfect competition and constant returns to scale.[21] It takes on the following Cobb-Douglas form:

$$y_j = A_j \left(\varsigma H_{Mj}\right)^{\alpha} \left(F_{Mj}\right)^{1-\alpha}, \tag{19}$$

where y_j is the output of the final good in workplace j, A_j is the final goods productivity, H_{Mj} is the total workplace employment, and F_{Mj} is the floor space used for production, which has price p_i. $\varsigma < 1$ is a scale parameter that accounts for the fact that workers consume on average less office space than residents consume in living space.[22]

With profit maximization, the total commercial floor space consumption at a certain location can be derived as follows:

$$F_{Mj} = \left(\frac{\omega_j^{1/\varepsilon}}{\alpha A_j}\right)^{1/(1-\alpha)} \varsigma H_{Mj}. \tag{20}$$

3.8 Production and amenity spillovers

We additionally assume that there are productivity spillovers from other locations. Consequently, the final good productivity, A_j, consists of an exogenous productivity part, $\bar{A}_j$, and an endogenous part defined by employment density in all other locations weighted by the travel time to get there:

$$A_j = \bar{A}_j \left(\sum_{i=1}^{I} \frac{H_{Mi}}{(1-\iota_i)L_i} \exp(-\delta \tau_{ij1})\right)^{\gamma}, \tag{21}$$

[21] We don't explicitly include *subsidised* employer car parking, as it is the consequence of a combination of distorting income taxation (employer parking is not taxed as a benefit in kind) and minimum parking requirements policies (Katz and Mankiw, 1985; Gutiérrez-i Puigarnau and Van Ommeren, 2011). In the Netherlands, minimum parking requirements are typically very mild.

[22] Following Koster (2024), ς is chosen such that the ratio of residential to commercial floor space use $(\sum_{j=1}^{I} F_{Rj})/(\sum_{j=1}^{I} F_{Mj})$ matches the observed total share in the data.

where δ captures the spatial decay of agglomeration, while γ captures the agglomeration elasticity. Because we find that the majority of business trips in the OViN data are made by car we assume that business trips are made by car.

Analogous to productivity, amenity spillovers travel along the road network and are generated through population density from locations around j weighted by the travel time to get there. Hence, amenities consist of exogenous component $\bar{B}_i$ and an endogenous component that is dependent on the population density at all other locations:

$$B_i = \bar{B}_i \left(\sum_{j=1}^{I} \frac{H_{Rj}}{(1-\iota_j)L_j} \exp(-\rho \tau_{ij1}) \right)^{\upsilon}, \tag{22}$$

where υ captures the amenity elasticity and ρ is the spatial decay of the amenities. To maintain consistency, we use car travel times here, as OViN surveys indicate that most shopping and recreational trips are made by car. However, a significant portion of shopping and recreational trips are also completed by bicycle. Therefore, in our structural estimation, we also present results where υ and ρ are estimated using bicycle travel times, τ_{ij2}, which leads to very similar results.

3.9 Land market and construction

The total floor space in each location, F_i, is equal to the sum of demand for residential and commercial floor space. Floor space is supplied in a competitive construction market and uses developed land $(1-\iota_i)L_i$ and capital K_i as input. We assume that capital has a common price r for all locations, while the price for floor space, p_i, varies between locations. We assume that the production function of floor space takes the Cobb-Douglas form

$$F_i = F_{Ri} + F_{Mi} = \Upsilon_i K_i^{\mu} \left((1-\iota_i) L_i \right)^{1-\mu}, \tag{23}$$

where Υ_i defines the innate supply conditions at location i, similar to Koster (2024). Given (23), the first-order condition for optimal capital use becomes

$$K_i^* = (\mu \Upsilon_i p_i)^{\frac{1}{1-\mu}} L_i \tag{24}$$

where we have normalised the price of capital, r, to 1.

3.10 Equilibrium

In equilibrium, profits and utility are maximised. Decisions on the improvement of the bicycle infrastructure network, which directly impact commuting costs, are taken by an absentee central planner and, hence, are not anticipated. An equilibrium requires the factor markets to clear. That means, given model parameters $\{\alpha, \beta, \mu, \gamma, \delta, \rho, \upsilon, \lambda, \eta, \varepsilon, \psi_2, \kappa_m, b_m\}$, $\forall m$, for any worker and land vectors with elements $\{H_{Ri}, H_{Mj}, F_{Ri}, F_{Mj}, R_{j1}\}$, any commuting time vector with elements $\{\tau_{ijm}\}$, $\forall m$, and any additional exogenous location characteristics with elements $\{B_i, C_i, \bar{A}_j, \bar{B}_i, K_i\}$, a competitive equilibrium is a set of endogenous vectors with elements $\{p_i, w_j, \theta_i\}$ that determine endogenously $\{\pi_{ijm}\}$ such that for all locations i the commuting market and the land market clear.[23]

Commuting-market clearing. In equilibrium, the number of workers in j equals the sum over the residential population that commutes to workplace j. Equivalently, the number of residents in i equals the sum over the workplace population that commutes from residence location i.

$$H_{Ri} = \sum_{j=1}^{I} \pi_{ij|j} H_{Mj} \quad \text{and} \quad H_{Mj} = \sum_{i=1}^{I} \pi_{ij|i} H_{Ri}. \tag{25}$$

Land-market clearing. An equilibrium requires the absence of arbitrage in the supply of residential and commercial floor space. Total demand for floor space should equal total supply as in equation (23). Residential land market clearing implies that the total supplied residential floor space (F_{Ri}) is demanded by residents in all locations i:

$$F_{Ri} = (1 - \theta_i) F_i = \frac{(1-\beta)\mathbb{E}[\bar{w}_i]}{p_i} H_{Ri}, \tag{26}$$

where $\mathbb{E}[\bar{w}_i] = \sum_{s=1}^{L} \pi_{is|i} w_s$ is the expected wage at the residence. Likewise, the commercial land market equilibrium requires that the total supplied commercial floor space (F_{Mj}) equals the demand for floor space by the production sector across all locations j:

$$F_{Mj} = \theta_j F_j = \left(\frac{\omega_j}{\alpha A_j C_j^{\frac{1}{\varepsilon}}} \right)^{\frac{1}{1-\alpha}} \varsigma H_{Mj} \tag{27}$$

[23]The equilibrium is unique as long as agglomeration forces are smaller than dispersion forces (see Allen and Arkolakis, 2014; Desmet et al., 2018; Allen et al., 2024). As we will show in the structural estimation later, agglomeration forces, both, in terms of productivity and amenity, play a minor role in our model in comparison to dispersion forces. We have confirmed this with a numerical analysis à la Allen et al. (2024) and can therefore conclude that the equilibrium is unique.

Welfare. The welfare effects of cycling and cycling infrastructure improvements will be evaluated by the change in expected utility once cycling is no longer an option or cycling infrastructure is removed. The change in expected utility can be interpreted as a Hicksian compensation variation. It indicates the additional income that is needed in order to reach the same level of utility than in the baseline scenario. The equivalent income that is needed to reach the initial utility level, $\Delta\bar{w}$, can be derived from the ratio of expected utilities using (8):

$$\Delta\bar{w} = \frac{\bar{u}_1}{\bar{u}_0} = \frac{\left(\sum_{r=1}^{I}\sum_{s=1}^{I} B_{i,1}C_{j,1}\left(\frac{w_{j,1}\exp\left(-\eta t_{ij,1}\right)}{p_{i,1}^{1-\beta}}\right)^{\varepsilon}\right)^{\frac{1}{\varepsilon}}}{\left(\sum_{r=1}^{I}\sum_{s=1}^{I} B_{i,0}C_{j,0}\left(\frac{w_{j,0}\exp\left(-\eta t_{ij,0}\right)}{p_{i,0}^{1-\beta}}\right)^{\varepsilon}\right)^{\frac{1}{\varepsilon}}} \tag{28}$$

where the subscripts $\{0, 1\}$ refer to the baseline and a counterfactual scenario, respectively.

4 Structural estimation

This section introduces the methodology to estimate the necessary model parameters to solve the general equilibrium model as presented in Section 3. Our model is estimated for the Netherlands for the baseline year 2016 and the 12,782 neighbourhoods. We exclude neighbourhoods that have neither workplace employment nor residents. The absence of economic activity is assumed to be due to exogenous location-specific fundamentals (*e.g.*, the presence of water) that do not change in the counterfactual analysis.

We set the expenditure of production on commercial floor space at $1 - \alpha = 0.15$ and the share of household expenditure on floor space at $1 - \beta = 0.31$, both in line with official figures from Statistics Netherlands (see Statistics Netherlands, 2016a,c, respectively). Furthermore, we assume the capital share in the construction costs $\mu = 0.75$ following Ahlfeldt et al. (2015) and Combes et al. (2010a). We set the commercial floor space scale parameter ς following the methodology in Koster (2024).

In the following sections, we detail the structural estimation of the remaining parameters, explaining the methodology and results. We begin by discussing the estimation of mode choice in Section 4.1 and the commuting cost parameters in Section 4.2. The congestion elasticity is estimated in Section 4.3, followed by the remaining parameters related to productivity and amenities in Section 4.4. Table 2 summarises all baseline parameters estimated and used in the model, serving as a useful reference for the subsequent discussion.

TABLE 2 – MODEL PARAMETER OVERVIEW

COBB-DOUGLAS PARAMETERS			
$1-\alpha$	0.150	Firm expenditure share on commercial floor space	Statistics Netherlands (2016a)
$1-\beta$	0.310	Worker expenditure share on residential floor space	Statistics Netherlands (2016c)
μ	0.750	Capital share in construction costs	Ahlfeldt et al. (2015) and Combes et al. (2010a)
MODE CHOICE PARAMETERS			
κ_1	-0.070	Commuting time cost, car	See column (6), Table 3
κ_2	-0.089	Commuting time cost, bicycle	See column (6), Table 3
κ_3	-0.148	Commuting time cost, walk	See column (6), Table 3
κ_4	-0.039	Commuting time cost, public transport	See column (6), Table 3
b_1	—	Preference shifter for car	Car is the reference mode
b_2	1.30	Preference shifter for bicycle	Calibrated, Statistics Netherlands data
b_3	1.97	Preference shifter for walk	Calibrated, Statistics Netherlands data
b_4	-2.19	Preference shifter for public transport	Calibrated, Statistics Netherlands data
ψ_2	0.972	Within correlation of non-car nest	See column (6), Table 3
GRAVITY PARAMETERS			
$\nu = \eta\varepsilon$	0.086	Commuting semi-elasticity	See column (1), Table 4
ε	6.03	Productivity heterogeneity	See column (1), Table 5
CONGESTION, AGGLOMERATION, AMENITIES			
λ	0.014	Congestion elasticity	See column (4), Table 6
λ_2	0.004	Congestion elasticity (squared term)	See column (4), Table 6
γ	0.023	Agglomeration elasticity	See column (2), Table 7
δ	0.078	Agglomeration decay	See column (2), Table 7
υ	0.045	Residential amenities elasticity	See column (4), Table 7
ρ	0.095	Residential amenity decay	See column (4), Table 7

Note: We report preference shifters b_m that are re-calibrated such that the mode shares according to the model equal the observed mode shares in the `OViN` survey for 2016-2022 (as in Table 1, Panel A), following the procedure outlined by Train (2009). Further details on the re-calibration exercise are provided in Section B.1 of the Appendix.

4.1 Mode choice estimation

Estimation and identification. The demand for travel modes can be expressed as the probability that a worker chooses mode m. Using the model structure, this probability can be decomposed into the probability that a worker chooses the nest containing m and the probability the worker chooses the mode from the mode options available in that nest:

$$\begin{aligned}\pi_{m|ij} &= \pi_{k|ij} \times \pi_{m|kij} \\ &= \frac{\exp\left(\psi_k t_{k|ij}\right)}{\sum_{n=1}^{2}\exp\left(\psi_n t_{n|ij}\right)} \times \frac{\exp\left(\frac{b_m}{\psi_k}+\frac{\kappa_m}{\psi_k}\tau_{ijm}\right)}{\sum_{n\in\mathbb{M}_k}\exp\left(\frac{b_n}{\psi_k}+\frac{\kappa_n}{\psi_k}\tau_{ijn}\right)}.\end{aligned} \tag{29}$$

One can estimate the demand for travel modes as defined by (29) using a Nested Logit model that is estimated with maximum likelihood. We employ trip-level data using the 2016-2022 OViN Survey (see Statistics Netherlands, 2016b). This Nested Logit approach will identify the sensitivity of commuting time for various modes, κ_m, and the within-correlation of mode choices of the non-car nest, ψ_2. This approach analyses mode choice conditional on the selected location pair, so that endogeneity concerns are limited. However, two potential issues merit attention: omitted variable bias and reverse causality.

Omitted variable bias could occur if workers with specific characteristics tend to choose certain modes. For example, low-income workers might predominantly bicycle or walk due to the lower cost, meaning that the estimated parameters may reflect not just sensitivity to commuting time, but also the behaviour of low-income workers. To address this, we control for various worker characteristics, including age, sex, education, and income, as well as day of the week (which accounts for day-of-the-week differences in travel times). Importantly, our analysis shows that excluding these controls has minimal impact on the results, suggesting that omitted variable bias because of sorting based on individual differences is not a major concern.

Reverse causality is another concern. For example, if many workers prefer to take public transport between locations i and j, local governments might invest in improving the public transport infrastructure at this link, thus reducing commute times. Another example is road congestion, where longer commuting times are the result of higher levels of demand. To address this issue, we apply a control function approach in which we instrument commuting times using the Euclidean distance between locations i and j, a measure that is not affected by the (mode-specific) flow of commuters between these locations. The control function approach implies that in the first stage of the estimation process, the mode-specific commuting times are regressed on the instrument with controls. The first stage errors are then included as control functions in the second-stage maximum likelihood estimation (Blundell and Powell, 2003).

Results. Table 3 presents the results. We start with a specification where we impose that the parameter κ is the same for all modes. We find a parameter of about -0.08, see column (1). This estimate remains virtually unchanged when we add worker characteristics (column (2)), or instrument for commuting time using the Euclidean distance (column (3)). This parameter is closely aligned with the findings of Ahlfeldt et al. (2015) and others.

In columns (4)-(6), we allow the κ to differ between modes.[24] Columns (4) and (5) suggest that there are non-negligible differences in the value of κ for different modes. This conclusion is reinforced when we instrument for commuting time in column (6), which is our preferred speci-

[24]To allow for this is nowadays common in the transport literature, see for example Kouwenhoven et al. (2014).

TABLE 3 – MODE CHOICE PARAMETER ESTIMATION

Dependent variable:	*Mode choice*					
	Single $\kappa_m = \kappa$			Mode-specific κ_m		
	(1)	(2)	(3)	(4)	(5)	(6)
	NL	NL	NL-CF	NL	NL	NL-CF
Commuting time, κ	-0.077***	-0.077***	-0.080***			
	(0.001)	(0.001)	(0.001)			
Car time, κ_1				-0.097***	-0.094***	-0.070***
				(0.002)	(0.002)	(0.004)
Bicycle time, κ_2				-0.093***	-0.093***	-0.089***
				(0.001)	(0.001)	(0.001)
Walk time, κ_3				-0.128***	-0.128***	-0.148***
				(0.005)	(0.005)	(0.006)
Public transport time, κ_4				-0.065***	-0.063***	-0.039***
				(0.002)	(0.002)	(0.003)
Non-car nest parameter, ψ_2	0.590***	0.603***	0.531***	0.849***	0.846***	0.972***
	(0.013)	(0.014)	(0.014)	(0.021)	(0.022)	(0.031)
Worker characteristics		✓	✓		✓	✓
First stage errors			✓			✓
Number of observations	130,240	130,240	130,240	130,240	130,240	130,240
Pseudo-R^2	0.293	0.295	0.304	0.307	0.317	0.320
First stage F-statistic			4,940			4,940

Notes: Nested logit model results. Data is sourced from the `OViN` survey for the years 2016-2022, where workers can choose between four modes (walk, bicycle, car, public transport). Worker characteristics refer to age, sex, education, income and day of the week. Car is the only mode within the car nest, so $\psi_1 = 1$. Columns (1) and (2) imposes the same κ for all modes. Columns (3) and (4) use mode-specific commuting time parameters κ_m. In columns (2) and (4), commuting time is instrumented with the Euclidean distance for each mode using a control function approach. Bootstrapped standard errors (100 replications) clustered at the i and j level in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

fication. We observe that the estimated parameters change when we instrument for commuting time, which did not occur when we assumed the same κ for all modes. This is reasonable because the causes of endogeneity differ between the various modes of transport.[25] In line with the idea that very few people bicycle or walk for very long times (*e.g.*, more than 60 minutes), and that for long distances, car drivers switch to train, we find (in absolute values): $\hat{\kappa}_4 < \hat{\kappa}_1 < \hat{\kappa}_2 < \hat{\kappa}_3$. Consequently, the marginal time costs are the highest for walking, followed by bicycling, car use, and public transport. These results are consistent with the notion that workers tend to walk shorter distances, then switch to cycling longer distances, then use the car, and then to using the train. The marginal costs of *distance* (rather than time) also depend on the speed of the specific mode. Typically, for car and public transport users, speed levels are much higher for commuting trips between cities compared to other modes, making the marginal cost of distance for car and public transport significantly lower than for cycling and walking.

[25] For instance, public transport infrastructure provision may be more responsive to commuting flows compared to the provision of sidewalks for pedestrians.

We also estimate the correlation parameter among the alternatives within the non-car nest, ψ_2. When κ is imposed to be the same for all modes, we get estimates within the range from 0.53 to 0.60, see columns (1)-(3). When we allow κ to be different among modes, ψ_2 is substantially closer to one, the value implied by the multinomial logit model. In our preferred specification in column (6), the estimate of ψ_2 is not statistically different from one.

Substitution between modes plays a central role in this paper. We aim to match the (aggregate) mode shares implied by the model with the ones we observe in the data. To achieve this, we recalibrate the mode-specific preference shifters b_m so that the mode shares are equal to the simulated mode shares using a recalibration procedure outlined by Train (2009) (we provide more details in Appendix B.1). The intercepts shown in Table 2 reveal that the workers prefer to cycle and walk over driving. In contrast, workers generally have a distaste for public transport compared to driving.

4.2 Gravity estimation, wages and commuting heterogeneity

Estimation and identification. To estimate the gravity equation, we assume that the generalised commuting time t_{ij} determines ij location choices of workers, similar to Tsivanidis (2023). Given the estimates of the mode choice parameters, the generalised commuting time t_{ij} can be calculated for each combination of residence and workplace. Given this generalised commuting time, we estimate the semi-elasticity of the commuting time ν.

To be more precise, first we calculate the nest-specific commuting times $t_{k=1|ij}$ and $t_{k=2|ij}$ using equation (14), given the estimated mode choice parameters $\{\hat{b}_m, \hat{\kappa}_m\}, \forall m$ and $\hat{\psi}_2$. Next, we calculate the generalised commute time t_{ij} using $\hat{t}_{k=1|ij}$ and $\hat{t}_{k=2|ij}$ and (5). Then, we log-linearise (3). Recall that $d_{ij} = \exp(\eta t_{ij})$. We then obtain:

$$\mathbb{E}[N_{ij}] = \exp(-\nu t_{ij} + \upsilon_i + \upsilon_j), \tag{30}$$

where $\nu = \eta\varepsilon$, $\mathbb{E}[N_{ij}] \equiv \pi_{ij}\bar{H}$ refers to the expected commuting flow, υ_j are residence and workplace fixed effects that absorb residential amenities, workplace amenities, and workplace wages. The commuting semi-elasticity ν is the parameter of interest. We estimate this parameter using (30) by PPML with observed commuting flow data from the 2016-2022 `OViN` Survey. Given 12,782 neighbourhoods, we observe flows between 163,379,524 flows neighbourhood pairs. PPML allows the observed commuting flow data to have a large share of zeros.

Given estimates of $\hat{\nu}$, we obtain the transformed wages $\omega_j = C_j w_j^{\varepsilon}$ using the commuting market clearing condition (25). Subsequently, we estimate the commuting heterogeneity parameter ε following the approach outlined in Ahlfeldt et al. (2015). The idea is that the variances

TABLE 4 – GRAVITY REGRESSION

Dependent variable:	*Commuting flow, N_{ij}*	
	(1) PPML	(2) PPML
Generalised commuting time, ν	0.086***	
	(0.003)	
Average commuting time		0.076***
		(0.000)
Residence fixed effects	✓	✓
Workplace fixed effects	✓	✓
Number of observations	163,379,524	163,379,524
Pseudo-R^2	0.45	0.37

Notes: Data from the 2016-2022 OViN Survey. Data refer to residence-work commuting flows N_{ij} and generalised commuting times t_{ij}, calculated using Table 3 and equation (5). Column (1) is a PPML regression with generalised commuting time. Column (2) is a PPML regression as in Ahlfeldt et al. (2015) without mode choice, where we use an average of the car and public transport commuting times as the commuting time between i and j. Bootstrapped standard errors (100 replications) clustered at the i and j levels in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

of observed log wages at a certain location, $\sigma^2_{\log w_j}$, and the variances of log transformed wages at this location, $\sigma^2_{\log \omega_j}$, are related as follows:

$$\sigma^2_{\log w_j} = \left(\frac{1}{\varepsilon}\right)^2 \sigma^2_{\log \omega_j} + \chi_j, \tag{31}$$

where χ_j refers to random error.

Results. Table 4 presents the results for the commuting time semi-elasticity. We find $\hat{\nu} = 0.086$. This estimate of the generalised commuting time decay is consistent with the estimate range found in previous studies that have estimated commuting time elasticities for one or more modes (Ahlfeldt et al., 2015; Dericks and Koster, 2021; Koster, 2024).

In column (2), we explore an alternative method to the above Nested Logit approach by calculating the average commuting times across car and public transport, following the method used in Ahlfeldt et al. (2015). This approach provides a very similar commuting time elasticity of 0.076, which increases confidence in the accuracy of the results.[26] The main disadvantage of using the average commuting time is that this method does not allow for any substitution between different modes. Consequently, our approach, which incorporates changes in mode choice in the counterfactual, is preferred, at least when these mode changes are non-negligible.

In Table 5 we present results from the estimation equation (31). First, using individual wage

[26]The latter elasticity is slightly lower. This makes sense as the average commuting time may be rewritten as the generalised commuting time plus measurement error. Random measurement error typically leads to a bias towards zero.

TABLE 5 – COMMUTING HETEROGENEITY PARAMETER

Dependent variable:	*Variance of observed wages*		
	Municipality level	District level	
	(1)	(2)	(3)
	OLS	OLS	OLS
Commuting heterogenity, $\hat{\varepsilon}$	6.03***	5.43***	5.76***
	(0.011)	(0.000)	(0.000)
Municipality fixed effects			✓
Number of observations	384	2,416	2,416
Pseudo-R^2	0.07	0.07	0.20

Notes: Wage data is sourced from the income registry data `SPOLIS` (*Salaris Polis Administratie*). Column (1) estimates equation (31) is based on variances at the municipality level. Columns (2) and (3) employ variances at the district level. Column (3) adds municipality fixed effects. Bootstrapped standard errors (100 replications) clustered at the i and j level in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

data, we calculate the wage variances at the municipality level (see column (1)). The implied $\hat{\varepsilon} = 6.03$ is very similar to previous estimates (Ahlfeldt et al., 2015, find a value of $\hat{\varepsilon} = 6.6$ for Berlin, while Koster, 2024, find $\hat{\varepsilon} = 5.5$ for England). Second, we calculate the wage variances at the district level, which leads to a very similar, but smaller, value of commuting heterogeneity (see column (2)). The latter result is not materially influenced when we additionally include municipality fixed effects, see column (3).

4.3 Car congestion parameters

Estimation and identification. We also estimate the congestion externality parameters, λ_1 and λ_2, based on the structure of the model. Rewriting (16) and taking logs gives the following equation:

$$\log\left(\frac{\tau_{ij1}}{\tau_{ij1}^{f}}\right) = \lambda_1 \Lambda_{Mj} + \lambda_2 (\Lambda_{Mj})^2 + \log(T_{Ri}) + \log(T_{Mj}) + \xi_{ij1}, \tag{32}$$

where T_{Ri} and T_{Mj} capture residence and workplace fixed-effects, respectively, and ξ_{ij1} is a residual that includes a common constant. Here, the dependent variable is the logarithm of the ratio of congested to free-flow car commuting costs, $\tau_{ij1}/\tau_{ij1}^{f}$, which is calculated using road network travel times for free-flow conditions and imputed `Google` travel times for congested conditions (see Appendix A.4 for details). The main independent variable of interest, car traffic density, Λ_{jM}, is equal to the ratio of the number of car commuters to the workplace location, H_{Mj1}, and the car lane kilometres per workplace location (see equation (17)), which is derived the `OSM` road network data.

There may be concern that when estimating equation (32) that Λ_{Mj} is endogenous due to reverse causality. For example, shorter commuting times could attract residents and workers

who prefer shorter commutes, increasing density. To address this, we employ a classical strategy used in the literature on agglomeration economies. More specifically, we instrument traffic densities using the population density of 1900 (see Gaigné et al., 2022). This historic instrument addresses the issue that decisions regarding workplace locations are influenced by the placement of infrastructure, as population density in 1900 is not expected to be correlated to current infrastructure supply, for example, because highways and roads suitable for cars did not exist around that time and commuting to work over longer distances was still uncommon.

This IV procedure implies that we estimate the congestion parameters of interest $\{\lambda_i\}$ in a two-step procedure. In the first step, we recover the workplace fixed effects, $\log \hat{T}_{Mj}$, with a regression where we condition only on residential fixed effects:

$$\log\left(\frac{\tau_{ij1}}{\tau^{f}_{ij1}}\right) = \log T_{Ri} + \log T_{Mj} + \tilde{\chi}_{ij1}.$$

Here, $\log T_{Ri}$ refers to a residential fixed effect, $\log T_{Mj}$ refers to a workplace location fixed effect and $\tilde{\chi}_{ij1}$ refers to an error term.

In the second step, we regress the recovered workplace fixed effects on traffic density, where we instrument the latter with the historic population density, $H_{Ri,1900}/((1-\iota)L_i)$, using a control function approach. The latter approach implies that one controls for the quadratic function of the first stage error, which is obtained by regressing Λ_{Mj} on the instrument. Consequently, we estimate the following equations:

$$\begin{aligned} \Lambda_{Mj} &= \tilde{\lambda}_1 \sum_i^I \exp(-\kappa_1 \tau_{ij1}) \frac{H_{Ri,1900}}{(1-\iota)L_i} + \tilde{\lambda}_2 \left(\sum_i^I \exp(-\kappa_1 \tau_{ij1}) \frac{H_{Ri,1900}}{(1-\iota)L_i}\right)^2 + \xi_j, \\ \log \hat{T}_{Mj} &= \lambda_1 \Lambda_{Mj} + \lambda_2 (\Lambda_{Mj})^2 + \beta_1 \hat{\xi}_j + \beta_2 (\hat{\xi}_j)^2 + \chi_j, \end{aligned} \tag{33}$$

where ξ_j and χ_j refer to error terms and β_1, β_2, $\tilde{\lambda}_1$, $\tilde{\lambda}_2$, λ_1 and λ_2 are coefficients to be estimated.

Results. Table 6 presents the results from estimating equations (32) using OLS and (33) using a control function approach where we use historic population density as an instrument. First, we represent a specification where we include car density but not its square. The OLS estimate, reported in column (1) suggests that the semi-elasticity of car density $\hat{\lambda}_1 = 0.015$. Accounting for endogeneity in column (2) increases it to 0.018, which implies a mild level of congestion. According to this specification, at the most congested locations (in Amsterdam), travel times are about 10% higher due to congestion. The size of this increase suggests marginal external costs that are considerably lower than observed in major cities around the world (Yang et al., 2020;

TABLE 6 – CONGESTION ELASTICITY

Dependent variable:	*Recovered workplace fixed effect,* $\log \hat{T}_{Mj}$			
	(1) OLS	(2) CF	(3) OLS	(4) CF
Car density, λ_1	0.013***	0.018***	0.013***	0.014***
	(0.000)	(0.001)	(0.000)	(0.000)
(Car density)^2, λ_2			0.004***	0.004***
			(0.000)	(0.000)
First stage error		✓		✓
Number of observations	12,782	12,782	12,782	12,782
First stage F-statistic		15,214		17,732

Notes: CF refers to control function approach. Column (1) estimates log workplace fixed effects on car density following (32), that has been divided in a two-step approach. In column (2), we instrument car density using historic population density in 1900. Column (3) estimates log workplace fixed effects on car density and car density squared. In column (4), we instrument car density with historic population density in 1900 and its squared counterpart. Bootstrapped standard errors (100 replications) clustered at the i and j levels in parentheses. *** p $<$ 0.01, ** p $<$ 0.05, * p $<$ 0.1. First stage results are reported in Table B2.

Russo et al., 2021), but at the same time confirms anecdotal evidence suggesting that traffic congestion in the Netherlands is mild.

Arguably, this specification is misspecified because it assumes that the semi-elasticity is a constant, and therefore does not allow for strongly convex effects as car traffic densities increase, for example, because congestion effects only become substantial once densities reach a certain threshold (Daganzo, 2007; Daganzo et al., 2011). Columns (3) and (4) therefore also include the square of car density. The OLS and the control function results support the view that the effect of car density is strongly convex. Both approaches indicate that $\hat{\lambda}_2 = 0.004 > 0$. These results imply that the semi-elasticity is close to zero for the lowest car densities and about 0.038 for the highest densities. Still, this specification also implies that congestion is very mild. In the most congested locations, travel times are still only 11% higher compared to free flow. We consider column (4) as the baseline specification.

4.4 Productivity and amenity elasticities

Estimation and identification. The final goods productivity, A_j, can be recovered from the model up to a normalisation constant. Given the optimal use of commercial floor space and the optimal use of labour, the final goods productivity can be written as:

$$A_j = \frac{1}{(1-\alpha)^{1-\alpha}\alpha^{\alpha}} p_j^{1-\alpha} \hat{\omega}_j^{\alpha/\hat{\varepsilon}}. \tag{34}$$

Analogously, the endogenous amenities, B_i, can be recovered given the optimal use of residential floor space and workers as follows:

$$B_i = \left(\frac{H_{Ri}}{W_i}\right)^{\hat{\varepsilon}} p_i^{1-\beta}, \tag{35}$$

where W_i is a measure of commuting market access, which can be expressed as:

$$W_i = \sum_{j=1}^{I} e^{-\hat{\nu} t_{ij}} \hat{\omega}_j. \tag{36}$$

The final goods productivity and endogenous amenities are then calculated using parameters of previous subsections $\{\alpha, \beta, \hat{\varepsilon}, \hat{\nu}\}$, estimated transformed wages $(\hat{\omega}_j)$, and observed data on housing prices (p_i). Log-linearising equations (21) and (22) gives

$$\log \hat{A}_j = \log(\bar{A}_j) + \gamma \log\left(\sum_{i=1}^{I} \frac{H_{Mi}}{(1-\iota_i)L_i} \exp(-\delta \tau_{ij1})\right) + \xi_{Mj} \tag{37}$$

$$\log \hat{B}_i = \log(\bar{B}_i) + \upsilon \log\left(\sum_{j=1}^{I} \frac{H_{Rj}}{(1-\iota_j)L_j} \exp(-\rho \tau_{ij1})\right) + \xi_{Ri}, \tag{38}$$

We estimate the parameters $\{\gamma, \delta\}$ and $\{\upsilon, \rho\}$ using non-linear least squares.

One potential concern with this approach is that spatially-weighted employment and residential densities may be correlated with unobserved location endowments. This implies that the most attractive locations could have higher densities for reasons not related to agglomeration economies, potentially leading to overestimation of γ and υ. To address this, we employ (again) population density data from 1900. This method, pioneered by Ciccone and Hall (1996) and widely applied in subsequent research (see Combes et al., 2010b), assumes that unobserved endowments are uncorrelated over such a long period. For instance, production in 1900 often required proximity to coal mines to minimize transport costs of inputs, but today, all such mines have been closed. We also incorporate travel-to-work-area fixed effects, which further bolster the exclusion restriction, as the correlation between unobserved location endowments over a period of more than 100 years *within* travel-to-work areas is likely negligible.[27] For readers still concerned that the exclusion restriction might not perfectly hold, we highlight that we will present sensitivity analyses showing that our counterfactual results remain robust even when assuming considerably different parameters for the agglomeration and amenity elasticities.[28]

[27] Travel-to-work-areas as defined by Statistic Netherlands are equal to European NUTS3 regions.

[28] These results are available upon request.

TABLE 7 – PRODUCTIVITY AND AMENITY ELASTICITIES

Dependent variable:	*Productivities,* $\log \hat{A}_j$		*Amenities,* $\log \hat{B}_i$			
	Car travel time		Car travel time		Bicycle travel time	
	(1) NLS	(2) NLS-CF	(3) NLS	(4) NLS-CF	(5) NLS	(6) NLS-CF
Agglomeration elasticity, γ	0.027***	0.023***				
	(0.001)	(0.001)				
Agglomeration decay, δ	-0.009***	0.078***				
	(0.001)	(0.001)				
Amenity elasticity, υ			0.037***	0.045***	0.078***	0.045***
			(0.004)	(0.006)	(0.003)	(0.006)
Amenity decay, ρ			0.007**	0.095***	0.378***	0.035***
			(0.002)	(0.010)	(0.024)	(0.003)
TTWA fixed effects	✓	✓	✓	✓	✓	✓
First stage error		✓		✓		✓
Number of observations	11,901	11,901	12,348	12,348	12,348	12,348
Pseudo R^2	0.82		0.15		0.31	
First stage F-statistic		5,672		296		358

Notes: NLS refers to non-linear least squares. CF stands for control-function approach. TTWA stands for a travel-to-work-area. Column (1) estimates equation (37) by regressing local productivities on the spatially weighted employment density around j. In column (2), we instrument for the spatially weighted employment density using population density in 1900. Column (3) estimates equation (38) by regressing local residential amenities on the spatially weighted residential density around i. In column (4), we instrument for the spatially weighted residential density using population density in 1900. In column (5) and (6), we repeat the estimation as in column (4) and (5) assuming that amenity spillovers are governed by cycling travel times, as compared to car travel times. Bootstrapped standard errors (100 replications) clustered at the i and j level in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

Results. Table 7 presents the results from the estimation of equations (37) and (38). Columns (1) and (2) focus on the productivity estimation, while columns (3)-(6) address the amenity estimation.

In column (1), we estimate an agglomeration elasticity of 0.027, which aligns well with the range suggested by the literature (Melo et al., 2009). We observe a very small, and even wrongly-signed, decay in agglomeration economies. However, when we turn to the instrumented estimates in column (2), we find a substantial decay parameter of $\hat{\delta} = 0.078$. This specification implies that the benefits of agglomeration decrease by about half after 10 minutes of driving. The estimated decay is somewhat less pronounced than in Ahlfeldt et al. (2015), who focused on a single city, but it is of a similar order of magnitude to Koster (2024), who also examined an entire country. The agglomeration elasticity ($\hat{\gamma} = 0.023$) remains largely unaffected by instrumenting for density.

In column (3), we estimate the amenity elasticity and its decay. The amenity elasticity is slightly larger now. Although the decay in amenities is statistically significant, it appears surprisingly flat, with most of the amenity effect remaining after only 20 minutes of travel. When we instrument the current residential density using the population density from 1900, the elasticity increases slightly to $\hat{\upsilon} = 0.045$ (see column (4)), indicating that the elasticity is

twice as large as for agglomeration. The decay in residential amenities increases substantially, $\hat{\rho} = 0.095$, with instrumentation strongly suggesting that amenities decay more strongly than productivity (see Ahlfeldt et al., 2015; Koster, 2024).

As a sensitivity analysis, in columns (5) and (6), we repeat the estimation allowing amenities to diffuse by bicycle travel times instead of car travel times. In the preferred specification in column (6), we find an amenity elasticity identical to that in the corresponding specification in column (4). Unsurprisingly, since cycling speeds are much lower than car speeds (14km/h versus 45km/h, see Table 1), the estimated amenity decay is less pronounced.

In the counterfactual analyses, we will use the estimate from columns (4) given that that the majority of recreational and shopping trips (51%) within the Netherlands are done by car (based on the `OViN` Survey).

5 Counterfactual analysis

5.1 Scenarios

In this section, we examine the impacts of cycling and cycling infrastructure on the urban spatial structure and welfare. Our analysis is structured around three distinct counterfactual scenarios. We refer to Appendix C for a detailed solution algorithm.

First, we simulate the *No cycling* scenario to assess how the spatial distribution of economic activity shifts when bicycling is not an option. This scenario represents a situation where workers are unable to choose cycling as a commuting option.[29] The road capacity previously allocated to bicycle lanes will now be reassigned to cars in the 40% most densely populated areas, where cycling and road infrastructure compete for scarce land. This *No cycling* scenario represents the most extreme case in our model and serves as a benchmark to highlight the significance of cycling.

Second, we focus on the provision of high-quality cycling infrastructure, specifically examining how the removal of separate cycleways and the subsequent expansion of car lanes in urban areas impacts the spatial economic structure. This is a relevant exercise because 70% of major roads have separate cycleways. Of those 70%, the majority (51%) cannot be used for cycling anymore implying that cyclists have to take a detour. Meanwhile, individuals have the same preference for cycling as in the baseline. We refer to this as the *No separate cycleways* scenario. Again, cycleways are substituted by road lanes in densely populated areas, analogously to the *No cycling*

[29]Technically, we modify the bicycle preference parameter so that $b_2 \to -\infty$, making bicycle commuting extremely undesirable. In addition, the mode shares in the corresponding baseline scenario are recalibrated to the observed mode shares reported in Table 1, by adjusting the mode choice preference shifters $b_1, .., b_4$. The change in the shifters is only slight. We will do a similar shifter adjustment in the baseline of other scenarios.

scenario. Consequently, the provision of separate cycleways affects the equilibrium through two channels: *(i)* an increased commuting time for cyclists, *(ii)* an increase in road lane capacity which affects car density and hence, congested commuting time.[30]

Third, the *Cycling speed increase* scenario analyses a hypothetical future situation in which bicycle commuting becomes faster due to the increased adaptation of electric bicycles and the provision of cycling highways. We assume that cycling becomes 10% faster (at negligible costs).

5.2 Aggregate effects

Main results. Table 8 outlines the the main results for the aggregate effects of the three counterfactual scenarios. The *No cycling* scenario (column 1, NC) is the most extreme, resulting in a non-negligible reduction in expected utility.[31] The shares for the alternative modes to bicycling increase strongly. Bicycle commuters predominantly shift to car commuting, which observes an increase of about 20 percentage points, with smaller shifts to walking – 3.1 percentage points – and public transport – 1.5 percentage points.[32]

These changes in travel behaviour go together with an increase in residential urban sprawl as indicated by a 1.8% decline of the Gini coefficient of spatial residential population. The deconcentration of workplace employment is less pronounced, as the Gini coefficient of spatial employment reduces by only -0.6%.[33] In line with that, the mean commuting distance rises substantially, by about 30%. The surge in car use induces households to move further from their workplaces, as predicted by the monocentric city model (Gin and Sonstelie, 1992; Glaeser et al., 2008). Consistent with car usage increases, average car speed drops by 2.2%. On average commuting time increases by about 14%, which is not negligible, but much less than the percentage increase in commuting distance.

The *No separate cycleways* scenario (column (2)), where separate cycleways are removed and cyclists must either share the road with cars or take detours, exhibits trends similar to the *No cycling* scenario. Although the effects are smaller, there is still a notable 1.2% decline

[30]The counterfactual scenario implies that we remove separate cycleways from the `OSM` network. For some roads, this means that bicycles can still use the road but share the road with cars, whereas other roads cannot be used anymore by cyclists. Therefore, we recalculate counterfactual bicycle commuting times between each neighbourhood pair. Furthermore, we calculate counterfactual road lane km for all neighbourhoods that increase within urban areas as a consequence of removing separate cycleways. This will endogenously affect the car density per neighbourhood and ultimately the commuting time by car. For more technical details on the implementation of that counterfactual scenario see Appendix D.1.

[31]The model suggests a reduction of 5.7%, which translates to an annual welfare loss of €18.5 billion, based on the 2016 median household income of €39,100.

[32]We keep public transport *supply* constant in this counterfactual, which may not be realistic. If the increase in demand leads to congestion in public transport (*e.g.*, limited seat availability or longer turnaround times at stops), it could result in further shifts toward car commuting and away from public transport, making our current estimates conservative.

[33]For reference, the baseline Gini of spatial employment is 0.745, the Gini coefficient spatial of residential population is 0.622, so substantially lower, consistent with that employment is more spatially concentrated.

TABLE 8 – AGGREGATE RESULTS OF COUNTERFACTUAL SCENARIOS

	Baseline estimation			Identical time costs $\kappa_m = \kappa$			No congestion $\lambda_1 = \lambda_2 = 0$		
Scenario:	**NC**	**NW**	**HS**	**NC**	**NW**	**HS**	**NC**	**NW**	**HS**
Δ Expected utility *(in %)*	-5.7	-1.2	0.8	-4.4	-1.5	0.9	-7.6	-1.8	1.0
Δ Spatial employment Gini *(in %)*	-0.6	-0.1	0.1	0.0	0.1	0.0	-0.7	-0.2	0.1
Δ Spatial residence Gini *(in %)*	-1.8	-0.5	0.2	-0.7	-0.1	0.0	-2.4	-0.4	0.3
Δ Total land rents *(in %)*	0.9	-0.2	0.0	-0.5	-0.3	0.1	-1.7	-0.1	-0.2
Δ Total output *(in %)*	1.4	-0.1	0.1	-0.1	-0.3	0.1	-0.2	-0.1	0.1
Δ Mean car speed *(in %)*	-2.2	0.2	0.4	-1.6	0.0	0.3	0.8	0.3	0.0
Δ Mean commuting time *(in %)*	14.0	1.2	0.0	8.2	1.3	-0.2	22.1	3.0	0.0
Δ Mean commuting distance *(in %)*	30.3	7.1	0.5	21.2	6.6	0.1	59.1	9.7	0.0
Δ Car share *(in %-points)*	20.4	4.5	-2.3	13.3	4.3	-2.0	20.9	3.6	-1.9
Δ Bicycle share *(in %-points)*	-25.0	-5.2	2.7	-25.0	-6.7	3.1	-25.0	-4.4	2.3
Δ Walk share *(in %-points)*	3.1	0.4	-0.2	6.8	1.0	-0.4	2.1	0.5	-0.2
Δ Public transport share *(in %-points)*	1.5	0.3	-0.2	4.9	1.4	-0.7	2.0	0.4	-0.2

Notes: **NC** = *No cycling*; **NW** = *No separate cycleways*; **HS** = *Cycling speed increase. Identical time costs* refers to a case where the commuting time cost parameter is not mode-specific but where we have identical time costs across modes (*i.e.*, a single κ , see Table 3, column 3). Mode shares in the baseline scenario of each panel are recalibrated and standardised to the observed mode shares reported in Table 1. In the *No congestion* sensitivity analysis, we use *free-flow* travel times instead of congested travel times and put the congestion parameters $\{\lambda_1, \lambda_2\}$ to zero.

in expected utility, which is equivalent to an annual welfare loss of approximately €3.5 billion. This can be considered a conservative estimate of the impact of not providing adequate cycling infrastructure, as we ignore the reduction in safety when cyclists have to share the road with cars. Given the substantial benefits of dedicated cycleways, it is clear that the advantages far outweigh the costs of building such infrastructure, because the costs of providing separate cycleways is below the costs of building roads.[34]

We now observe rather moderate increases in car use – by 4.5 percentage points. Interestingly, car speeds slightly increase. The latter would not be possible in a transport model where commuting distances are given, but is consistent with that the removal of separate cycleways induces households to relocate away from city centres to the suburbs, where car speeds are typically higher. In line with this, commute distances increase moderately by about 7%.

In the final scenario, we explore a 10% increase in cycling speeds, e.g. due to the adoption of electric bicycles, leading to a non-negligible 0.8% rise in expected utility. The higher cycling speeds result in an increase of 2.8 percentage points in the share of commuters using bicycles, indicating that higher speeds nontrivially enhance the appeal of cycling. Nevertheless, car speed remains rather constant and although cities become slightly more compact, in essence the spatial

[34]For instance, the average costs of the car infrastructure is projected to be 47 euros per 1,000 passenger km, while for cycling this number is only 22 (Schroten et al., 2022).

structure hardly changes.

Sensitivity analysis: identical time costs. In contrast to the transport economics literature, it is common in the quantitative spatial model literature to assume that κ is the same for all travel modes, implying that the marginal time costs are the same for all travel modes (Tsivanidis, 2023). To assess whether allowing for mode-specific commuting costs is important, we re-estimate our counterfactuals based on estimates reported in column (3) of Table 3, which impose that the marginal time costs are the same for all travel modes, i.e. $\kappa_m = \kappa$.

Table 8 reveals important differences across all three scenarios. While commuting distances (and times) still increase under the assumption that marginal time costs are the same across all modes, the increase is now much smaller. Consequently, changes in the urban spatial structure are also more moderate when marginal time costs are assumed to be identical, as reflected by smaller changes in the reported Gini coefficients. These findings are consistent with that in the *No cycling* and *No separate cycleways* scenarios, the shift to car use is systematically less when imposing the identical time cost restriction. For instance, in the *No cycling* scenario, the car share increases by just 13 percentage points, as opposed to 20 percentage points we have reported above. These findings suggest that when data availability allows, future research in welfare effects from infrastructure investments should consider mode-specific time costs in the analysis.

Sensitivity analysis: no congestion. In the last three columns of Table 8 we have investigated the role of car congestion. We recalibrated the model under the assumption that all roads are uncongested. By so doing we assume car speeds to be equal to their free-flow levels. Although the external costs of car congestion in Dutch cities tend to be low compared to other countries, disregarding car congestion still has notable effects, and exactly in the direction predicted by the monocentric city model if cyclists predominantly move to car use (see our discussion of Figure 2): changes in spatial structure become more pronounced when cycling is reduced, and commuting distances (and times) increase more significantly. This may appear counterintuitive, but it is important to recognise that reducing congestion reduces the marginal costs of car travel. As a result, more people will switch to driving and drive longer distances.[35]

[35]We also observe a much larger drop in expected utility in the *No cycling* scenario. While this makes sense, it is somewhat artificial, as in a world without congestion costs, the preference for cycling would need to be much stronger to explain the current level of cycling. As a result, removing cycling would lead to a more significant welfare loss.

5.3 Local spatial differences

Here we consider the local spatial differences in the effects for the three scenarios. Let us start with the *No cycling* scenario for which we show the changes in residential density in Figure 5a. In line with the reduction in the Gini coefficient observed earlier, residential densities decrease significantly in inner cities, especially around the 6 largest cities, which are labelled in our figures. Densities in inner-city areas drop by about 14%. These findings suggest that the modal shift towards cars, driven by the reduction in cycling, leads to increased urban sprawl and lower average population densities, which aligns well with reduced-form empirical evidence (Glaeser and Kahn, 2004; Ostermeijer et al., 2022). Workplace densities also decline, although to a lesser extent than residential densities. Consistent with this, floor space prices fall in inner cities due to reduced demand, while prices increase in suburban and peripheral areas (see Appendix D.2).

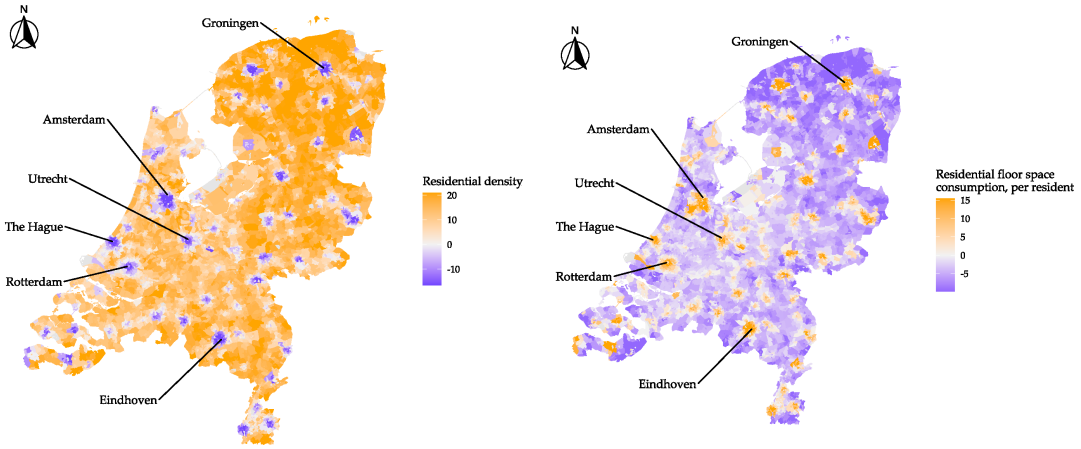

(a) Residential density Δ%

(b) Residential floor space consumption Δ%

Notes: This figure illustrates the local spatial differences of the *No cycling* scenario with respect to the model baseline. The left figure (a) shows changes in residential density (in percent), while the right figure (b) shows changes of residential floor space consumption (in percent). Percentages displayed are held between the 5^{th} and 95^{th} percentiles. Both figures display the names and locations of the six largest cities in the Netherlands.

Figure 5 – Local spatial differences in the *No cycling* scenario

Figure 5b highlights the changes in residential floor space consumption per person in different areas. As floor space prices drop in inner cities, consumption in these areas increases by at least 10%. In contrast, in suburban and peripheral regions, higher floor space prices lead to a reduction in floor space consumption.

We find that the wage differences in the workplace in the counterfactual scenarios are minimal. In the *No cycling* scenario, the average wage change is only −0.48% nationwide and 0.23% for the 6 largest cities. This result aligns with the relatively limited spatial variation in wage levels

in the Netherlands (see Groot et al., 2014).

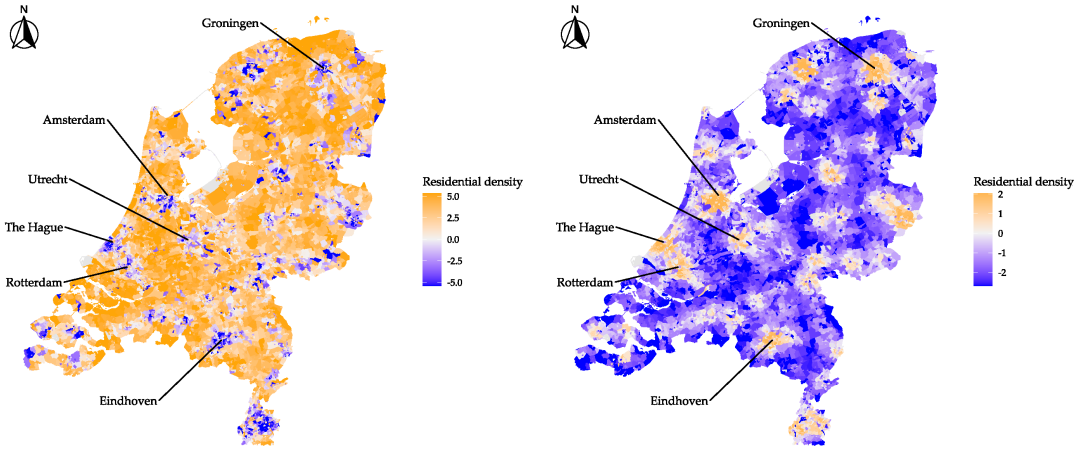

(a) Residential density Δ% (b) Residential density Δ%

Notes: This figure illustrates the local spatial differences arising from analysing the *No cycleways* scenario (left figure, (a)) and the *Cycling speed increase* scenario (right figure, (b)). Both figures show changes in residential density (in percent). Percentages displayed are held between the 5^{th} and 95^{th} percentile. Both figures display the names and locations of the six largest cities in the Netherlands.

Figure 6 – Local spatial differences in the *No cycleways* (a) and *Cycling speed increase* (b) scenarios

In Figure 6a, we analyse the effect on residential density under the *No separate cycleways* scenario. We find that replacing dedicated cycleways with road lanes within cities leads to suburbanisation, although to a lesser extent than in the *No cycling* scenario. Without separate cycleways, suburban growth reaches up to 5%, compared to 20% in the *No cycling* scenario. As detailed in Appendix D.2, the density of the workplace and the prices of floor space exhibit qualitative changes similar to those observed in the *No cycling* scenario, but with a smaller overall impact.

By contrast, making cycling more attractive by increasing speeds has a positive impact on inner-city density. As depicted in Figure 6b, a 10% rise in cycling speeds leads to an increase in inner-city densities of up to 2%. Workplace densities follow a similar trend but the impacts are smaller. Floor space prices also see a modest rise in cities, although only by about 0.5% (see again Appendix D.2). Interestingly, in certain rural areas in the north of the Netherlands, prices also experience a slight increase, likely due to the limited availability of public transport, making improvements in cycling speed particularly beneficial.

6 Conclusion

This paper explores the impact of cycling and cycling infrastructure on urban spatial structure and overall welfare. Our key methodological contribution lies in incorporating mode choice and including cycling into a spatial general equilibrium model. This approach advances recent work in the field (see, *e.g.*, Severen, 2023; Tsivanidis, 2023; Koster, 2024). In our model, the choice to cycle versus other commuting options (such as driving) interacts with residential and workplace location decisions, traffic congestion, and the resulting commuting times and distances. Importantly, we allow for mode-specific commuting time costs and correlation between various transport modes. This setup then allows us to carefully examine the trade-offs between cycling and car commuting, which are especially relevant in densely populated areas. Additionally, our use of fine-grained spatial data is a significant contribution, as cycling plays a crucial role in short-distance commutes, particularly those up to 5 km.

We estimate the model for the Netherlands, where cycling for commuting is very popular with cycling commuting shares of about 25%. However, the share of cycling is about 55% for shorter trips within cities, highlighting that mode choice highly depends on spatial differences in residential (and employment) densities. Using the structure of the model and the estimated model parameters, we conduct a series of counterfactual experiments in order to evaluate how changes in bicycle commuting and infrastructure impact the urban spatial structure, as captured by residential density, commuting distances, times, and mode shares.

We demonstrate that eliminating cycling leads to a 20 percentage-point increase in car use, which subsequently causes a 30% rise in commuting distance, a 14% increase in commuting time, and a 2% decline in car speed due to increased traffic congestion. As a result, the lower per-kilometre costs of driving, compared to cycling, contribute to greater urban sprawl. Overall, eliminating bicycle commuting reduces worker welfare by about 6%. These findings confirm that cycling not only improves welfare but also helps maintain more compact cities. Further, removing separate bicycle lanes moderately increases car use and induces moderate increases in urban sprawl – commuting distances increase by 6% – but still sizable reductions in welfare (about 1.2%). The increase in cycling speed, anticipated due to the increased adoption of electric bicycles, increases the use of bicycles, but has a limited impact on the spatial structure of the economy.

Our results are consistent with the monocentric city model, which predicts that individuals using modes with higher marginal costs prefer to live closer to the city center (Gin and Sonstelie, 1992; Glaeser, 2008). This implies that eliminating cycling increases urban sprawl. In line with these predictions, we also find that lower levels of car congestion amplify the reported effects on

spatial structure.

These findings highlight the critical role of maintaining and enhancing cycling infrastructure in shaping future urban development, alleviating congestion, and boosting overall welfare. The results offer compelling evidence for policymakers to prioritise cycling infrastructure as a cornerstone of sustainable urban transport planning. However, it is important to note that our counterfactual scenarios were carried out in a country with a strong cultural preference for cycling. As such, the external validity of our results is likely limited to countries with similar historic cycling preferences but currently less developed cycling infrastructure, such as Germany, England, Belgium, France, China and Japan. At the same time, our model could also be useful for other countries, as long as the mode-specific preference parameters are adjusted so that they reflect the prevailing mode shares.

References

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

Akar, G. and Clifton, K. J. (2009). Influence of individual perceptions and bicycle infrastructure on decision to bike. *Transportation Research Record: Journal of the Transportation Research Board*, (2140):165–172.

Akbar, P., Couture, V., Duranton, G., and Storeygard, A. (2023). Mobility and congestion in urban india. *American Economic Review*, 113(4):1083–1111.

Allam, Z., Bibri, S. E., Chabaud, D., and Moreno, C. (2022). The '15-minute city' concept can shape a net-zero urban future. *Humanities and Social Sciences Communications*, 9(1):1–5.

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

Allen, T. and Arkolakis, C. (2022). The Welfare Effects of Transportation Infrastructure Improvements. *The Review of Economic Studies*, 89(6):2911–2957.

Allen, T., Arkolakis, C., and Li, X. (2024). On the equilibrium properties of spatial models. *American Economic Review: Insights*, 6(4):472–89.

ANWB (1900). *De Kampioen*. Google Books.

ANWB (2024). De aanleg van de nederlandse fietspaden. Accessed: 2024-08-06.

Arnott, R. (2013). A bathtub model of downtown traffic congestion. *Journal of Urban Economics*, 76:110–121.

Aultman-Hall, L., Hall, F. L., and Baetz, B. B. (1998). Analysis of bicycle commuter routes using geographic information systems: Implications for bicycle planning. *Transportation Research Record: Journal of the Transportation Research Board*, (1578):102–110.

Baum-Snow, N. (2007). Did highways cause suburbanization? *The quarterly journal of economics*, 122(2):775–805.

Beenackers, M., Foster, S., Kamphuis, C., Titze, S., Divitini, M., Knuiman, M., and Giles-Corti, B. (2012). Taking up cycling after residential relocation built environment factors. *American Journal of Preventive Medicine*, 42(6):610–615.

Blundell, R. and Powell, J. (2003). Endogeneity in Nonparametric and Semiparametric Regression Models. In Dewatripont, M., Hansen, L., and Turnovsky, S., editors, *Advances in Economics and Econometrics: Theory and Applications*. Cambridge University Press, Cambridge.

Broach, J., Dill, J., and Gliebe, J. (2012). Where do cyclists ride? a route choice model developed with revealed preference gps data. *Transportation Research Part A: Policy and Practice*, 46(10):1730–1740.

Bruno, M., Dekker, H.-J., and Lemos, L. L. (2021). Mobility protests in the netherlands of the 1970s: Activism, innovation, and transitions. *Environmental Innovation and Societal Transitions*, 40:521–535.

Buehler, R. (2012). Determinants of bicycle commuting in the washington, dc region: The role of bicycle parking, cyclist showers, and free car parking at work. *Transportation Research Part D: Transport and Environment*, 17(7):525–531.

Buehler, R. and Dill, J. (2016). Bikeway networks: A review of effects on cycling. *Transport reviews*, 36(1):9–27.

Cervero, R. and Duncan, M. (2003). Walking, bicycling, and urban landscapes: Evidence from the san francisco bay area. *American Journal of Public Health*, 93(9):1478–1483.

Cervero, R., Sarmiento, O. L., Jacoby, E., Gomez, L. F., and Neiman, A. (2009). Influences of built environments on walking and cycling: Lessons from bogota. *International Journal of Sustainable Transportation*, 3(4):203–226.

Ciccone, A. and Hall, R. E. (1996). Productivity and the density of economic activity. *The American Economic Review*, pages 54–70.

Combes, P.-P., Duranton, G., Gobillon, L., and Roux, S. (2010a). Estimating Agglomeration Economies with History, Geology, and Worker Effects. In *Agglomeration Economics*, pages 15–66. University of Chicago Press.

Combes, P.-P., Duranton, G., Gobillon, L., and Roux, S. (2010b). Estimating agglomeration economies with history, geology, and worker effects. In Glaeser, E. L., editor, *Agglomeration economics*, pages 15–66. University of Chicago Press.

CROW (2017). *Design Manual for Bicycle Traffic*. REC28. CROW, The Netherlands. Productgroep: Drukwerk / Publicatie, Soort: Eenmalige aanschaf, Datum uitgifte: 01-01-2017, Aantal pagina's: 300, Druk: 0.

Daganzo, C. F. (2007). Urban gridlock: macroscopic modeling and mitigation approaches. *Transportation Research Part B: Methodological*, 41(1):49–62.

Daganzo, C. F., Gayah, V. V., and Gonzales, E. J. (2011). Macroscopic relations of urban traffic variables: bifurcations, multivaluedness and instability. *Transportation Research Part B: Methodological*, 45(1):278–288.

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

Desmet, K., Nagy, D., and Rossi-Hansberg, E. (2018). The geography of development. *Journal of Political Economy*, 126(3):903 – 983.

Dill, J. (2009). Bicycling for transportation and health: The role of infrastructure. *Journal of Public Health Policy*, 30:S95–S110.

Dingel, J. I. and Tintelnot, F. (2020). Spatial economics for granular settings. Working Paper 27287, National Bureau of Economic Research.

Eaton, J. and Kortum, S. (2002). Technology, Geography, and Trade. *Econometrica*, 70(5):1741–1779. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-0262.00352.

ECF (2023). Ratio of cycle tracks to main roads (plus information on surfaces).

Fosgerau, M. (2015). Congestion in the bathtub. *Economics of Transportation*, 4(4):241–255.

Gaigné, C., Koster, H. R. A., Moizeau, F., and Thisse, J.-F. (2022). Who lives where in the city? Amenities, commuting and income sorting. *Journal of Urban Economics*, 128:103394.

Gin, A. and Sonstelie, J. (1992). The streetcar and residential location in nineteenth century philadelphia. *Journal of Urban Economics*, 32(1):92–107.

Glaeser, E. L. (2008). *Cities, agglomeration, and spatial equilibrium*. OUP Oxford.

Glaeser, E. L. and Kahn, M. E. (2004). Sprawl and urban growth. In *Handbook of regional and urban economics*, volume 4, pages 2481–2527. Elsevier.

Glaeser, E. L., Kahn, M. E., and Rappaport, J. (2008). Why do the poor live in cities? the role of public transportation. *Journal of urban Economics*, 63(1):1–24.

Groot, S. P., de Groot, H. L., and Smit, M. J. (2014). Regional wage differences in the netherlands: Micro evidence on agglomeration externalities. *Journal of Regional Science*, 54(3):503–523.

Gutiérrez-i Puigarnau, E. and Van Ommeren, J. N. (2011). Welfare effects of distortionary fringe benefits taxation: The case of employer-provided cars. *International Economic Review*, 52(4):1105–1122.

Handy, S., Boarnet, M. G., Ewing, R., and Killingsworth, R. E. (2002). How the built environment affects physical activity: Views from urban planning. *American Journal of Preventive Medicine*, 23(2):64–73.

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

Heinen, E., Van Wee, B., and Maat, K. (2010). Commuting by bicycle: an overview of the literature. *Transport reviews*, 30(1):59–96.

Katz, A. and Mankiw, N. G. (1985). How should fringe benefits be taxed? *National Tax Journal*, 38(1):37–46.

Knol, W. C., Kramer, H., and Gijsbertse, H. (2004). *Historisch grondgebruik Nederland: Een landelijke reconstructie van het grondgebruik rond 1900*. Alterra, Wageningen.

Koster, H. R. A. (2024). The welfare effects of greenbelt policy. *Economic Journal*, 134(657):363–401.

Koster, H. R. A., Hayakawa, K., Tabuchi, T., and Thisse, J.-F. (2024). High-speed rail and the spatial distribution of economic activity: Evidence from Japan's shinkansen. *CEPR Discssion Paper*.

Kouwenhoven, M., de Jong, G. C., Koster, P., van den Berg, V. A., Verhoef, E. T., Bates, J., and Warffemius, P. M. (2014). New values of time and reliability in passenger transport in the netherlands. *Research in Transportation Economics*, 47:37–49.

Krizek, K. J. and Johnson, P. J. (2006). Proximity to trails and retail: Effects on urban cycling and walking. *Journal of the American Planning Association*, 72(1):33–42.

Loomis, J. (2011). What's to know about hypothetical bias in stated preference valuation studies? *Journal of Economic Surveys*, 25(2):363–370.

Martens, K. (2004). The bicycle as a feedering mode: experiences from three european countries. *Transportation Research Part D: Transport and Environment*, 9(4):281–294.

McFadden, D. (1978). Modelling the Choice of Residential Location. *Cowles Foundation Discussion Papers*.

Melo, P. C., Graham, D. J., and Noland, R. B. (2009). A meta-analysis of estimates of urban agglomeration economies. *Regional science and urban Economics*, 39(3):332–342.

Miyauchi, Y., Nakajima, K., and Redding, S. (2021). The Economics of Spatial Mobility: Theory and Evidence Using Smartphone Data. Technical Report w28497, National Bureau of Economic Research, Cambridge, MA.

Monte, F., Redding, S. J., and Rossi-Hansberg, E. (2018). Commuting, Migration, and Local Employment Elasticities. *American Economic Review*, 108(12):3855–3890.

Moreno, C., Allam, Z., Chabaud, D., Gall, C., and Pratlong, F. (2021). Introducing the "15-minute city": Sustainability, resilience and place identity in future post-pandemic cities. *Smart cities*, 4(1):93–111.

Nelson, A. C. and Allen, D. (1997). If you build them, commuters will use them: Association between bicycle facilities and bicycle commuting. *Transportation Research Record: Journal of the Transportation Research Board*, (1578):79–83.

Oldenziel, R. and Albert De La Bruhèze, A. (2011). Contested Spaces. *Transfers*, 1(2):29–49.

Ostermeijer, F., Koster, H. R. A., Van Ommeren, J., and Nielsen, V. M. (2022). Automobiles and urban density. *Journal of Economic Geography*, 22(5):1073–1095.

Padgham, M. (2019). dodgr: An R Package for Network Flow Aggregation. *Findings*. Publisher: Findings Press.

Pucher, J. and Buehler, R. (2006). Why canadians cycle more than americans: A comparative analysis of bicycling trends and policies. *Transport Policy*, 13(3):265–279.

Pucher, J. and Dijkstra, L. (2003). Promoting safe walking and cycling to improve public health: Lessons from the netherlands and germany. *American Journal of Public Health*, 93(9):1509–1516.

Pucher, J., Dill, J., and Handy, S. (2010). Infrastructure, programs, and policies to increase bicycling: An international review. *Preventive Medicine*, 50(Suppl 1):S106–S125.

Ranking Royals (2023). Top 10 countries with the highest bicycle usage. Accessed: 2024-08-05.

Rayaprolu, H. S., Llorca, C., and Moeckel, R. (2020). Impact of bicycle highways on commuter mode choice: A scenario analysis. *Environment and Planning B: Urban Analytics and City Science*, 47(4):662–677.

Redding, S. J. (2023). Quantitative Urban Models: From Theory to Data. *Journal of Economic Perspectives*, 37(2):75–98.

Rietveld, P. and Daniel, V. (2004). Determinants of bicycle use: do municipal policies matter? *Transportation research part A: policy and practice*, 38(7):531–550.

Russo, A., Adler, M. W., Liberini, F., and van Ommeren, J. N. (2021). Welfare losses of road congestion: Evidence from rome. *Regional Science and Urban Economics*, 89:103692.

Saelens, B. E., Sallis, J. F., and Frank, L. D. (2003). Environmental correlates of walking and cycling: Findings from the transportation, urban design, and planning literatures. *Annals of Behavioral Medicine*, 25(2):80–91.

Schoner, J. E. and Levinson, D. M. (2014). The missing link: Bicycle infrastructure networks and ridership in 74 us cities. *Transportation*, 41(6):1187–1204.

Schroten, A., Leestemaker, L., and Scholten, P. (2022). De prijs van een reis. Technical report, CE Delft.

Severen, C. (2023). Commuting, Labor, and Housing Market Effects of Mass Transportation: Welfare and Identification. *The Review of Economics and Statistics*, 105(5):1073–1091.

Small, K. A., Verhoef, E. T., and Verhoef, E. (2007). *The economics of urban transportation*. Routledge, London.

Statistics Netherlands (2016a). Bedrijfsleven; arbeids- en financiële gegevens. Statistics Netherlands.

Statistics Netherlands (2016b). OViN Survey. Survey data. Statistics Netherlands.

Statistics Netherlands (2016c). Woonlasten huishoudens. Statistics Netherlands.

Teulings, C. N., Ossokina, I. V., and de Groot, H. L. (2018). Land use, worker heterogeneity and welfare benefits of public goods. *Journal of Urban Economics*, 103:67–82.

Thomas, B. and DeRobertis, M. (2013). The safety of urban cycle tracks: A review of the literature. *Accident Analysis & Prevention*, 52:219–227.

Titze, S., Stronegger, W. J., Janschitz, S., and Oja, P. (2008). Association of built-environment, social-environment and personal factors with bicycling as a mode of transportation among austrian city dwellers. *Preventive Medicine*, 47(3):252–259.

Ton, D., Bekhor, S., Cats, O., Duives, D. C., Hoogendoorn-Lanser, S., and Hoogendoorn, S. P. (2020). The experienced mode choice set and its determinants: Commuting trips in the netherlands. *Transportation Research Part A: Policy and Practice*, 132:744–758.

Ton, D., Duives, D. C., Cats, O., Hoogendoorn-Lanser, S., and Hoogendoorn, S. P. (2019). Cycling or walking? Determinants of mode choice in the Netherlands. *Transportation Research Part A: Policy and Practice*, 123:7–23.

Train, K. (2009). Discrete Choice Methods with Simulation.

Tsivanidis, N. (2023). Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio.

Winters, M., Davidson, G., Kao, D., and Teschke, K. (2011). Motivators and deterrents of bicycling: Comparing influences on decisions to ride. *Transportation*, 38(1):153–168.

Xing, Y., Handy, S. L., and Mokhtarian, P. L. (2010). Factors associated with proportions and miles of bicycling for transportation and recreation in six small us cities. *Transportation Research Part D: Transport and Environment*, 15(2):73–81.

Yang, J., Purevjav, A.-O., and Li, S. (2020). The marginal cost of traffic congestion and road pricing: evidence from a natural experiment in beijing. *American Economic Journal: Economic Policy*, 12(1):418–453.

Yin, A., Chen, X., Behrendt, F., Morris, A., and Liu, X. (2024). How electric bikes reduce car use: A dual-mode ownership perspective. *Transportation Research Part D: Transport and Environment*, 133:104304.

Cycling cities: Mode choice, car congestion, and urban structure

Online Appendix

Tijl Hendrich[i] Hans R.A. Koster[ii] Nicole Loumeau[iii] Jos van Ommeren[iv]

Abstract – In this Appendix, we provide more information on the data used, as well as additional description and results regarding the structural estimation and counterfactual analysis. In Appendix A we offer more information on the data sources and definitions. Appendix B focuses on additional information regarding the structural estimation, including how we calibrated the aggregate mode shares to observed data, various first stage results of key parameter estimations and a description of how we imputed mode shares and congested travel times that were not observed in the data. Appendix C provides the solution algorithm of the counterfactual analysis, while Appendix D shows various additional maps on the counterfactual analysis.

Keywords – Cycling, bicycles, car congestion, agglomeration, density, infrastructure, commuting, 15-minute cities

JEL codes – R11, R12, R13, R14.

[i]CPB Netherlands Bureau for Economic Policy Analysis, corresponding author, email: t.j.m.hendrich@cpb.nl; Department of Spatial Economics, Vrije Universiteit Amsterdam.

[ii]Department of Spatial Economics, Vrije Universiteit Amsterdam, email: h.koster@vu.nl; Tinbergen Institute, CEPR.

[iii]CPB, Netherlands Bureau for Economic Policy Analysis, email: n.loumeau@cpb.nl.

[iv]Department of Spatial Economics, Vrije Universiteit Amsterdam, email: jos.van.ommeren@vu.nl; Tinbergen Institute.

A Data sources and definitions

A.1 Neighbourhoods

Our main data source is the universe of *non-public* Dutch 2016 microdata compiled by Statistics Netherlands. All data is based on administrative sources unless indicated otherwise. Our unit of analysis is the neighbourhood (or the 'buurt' in Dutch) as defined by Statistics Netherlands. In 2016, we have information on 12,782 neighbourhoods, containing on average only about 600 households. We exclude a small number of neighborhoods located on the Wadden Islands due to the lack of available data on ferry travel times. Table A1 gives an overview of all relevant neighbourhood statistics, for which we explain the source in more detail below.

Productivity and

A.2 Population

Wages. For all Dutch workers, we obtain data on monthly income, hours worked and their employer in our base year 2016 from the income registry data `SPOLIS` (*Salaris Polis Administratie*), which we rework to the equivalent hourly wages. We discard workers that receive unemployment and disability benefits, state pensions, firm owners and temporary and payroll workers and the self-employed, since they either do not have a fixed work location or their travel behaviour is likely to be governed by factors outside our model. We also drop workers that work less than 8 hours per week since they are unlikely to commute to work multiple times per week on a regular basis. This way, we end up with a total number of observations of $N \approx 5,000,000$.

Table A1 – Summary statistics at the neighbourhood level

	Mean	Std. Dev.	10^{th}	50^{th}	90^{th}
Residence population	398	546	10	192	1045
Working population	398	975	4	106	977
Population in 1900	400	1093	0	124	898
Median wage per hour (in euro)	19	4	13	19	24
Floor price per m^2 (in euro)	1276	440	842	1197	1772
neighbourhood area (in km^2)	2.6	5.3	0.2	0.7	7.4
Road length (in km)	2.1	5.3	0	0.3	5.5

Notes: The data are sourced from the respective datasets mentioned in subsection A. Road length corresponds to the road capacity used for calculating traffic density in our model (see equation (17)). It refers to the number of lane km per neighbourhood for all primary and secondary roads. Road lane km of tertiary and residential roads are not included.

Residences. To obtain the neighbourhood of residence of each person living in the Netherlands, we match records from the municipal population register (`GBA`) to the register of all Dutch addresses (`GBAADRESOBJECTBUS`) on 6-digit zip code level where each zip code represents about 30 houses. Subsequently, we map the zip codes to the set of 12,782 neighbourhoods, which we consider in our analysis. We perform this procedure for the model base year 2016. The procedure covers the complete Dutch population. Finally, we merge the residences data to the wages dataset to pair individual workers and their residences to their respective employers.

Workplaces. On the workplace side, we extract all firms and all of their branch locations from the national firm register (`ABR`) and match them with the neighbourhoods they are located in. Since the wages data only contains the employer, but not the specific branch where a given employee works, we set the workplace to the location closest the worker residence if a firm has multiple branches, as in Gaigné et al. (2022). We now have constructed a dataset of about 5 million individual Dutch workers coupled to their residence and work neighbourhoods.

Historic population. We use historic population data from 1900 as an instrumental variable to estimate the congestion elasticity and parameters on agglomeration economies pertaining to productivity and amenities. This population data comes from *NLGIS* for municipalities in 1900, which were smaller and comparable to the size of a large neighbourhood today. The local population distribution is imputed by mapping buildings and assuming a uniform population per building within each municipality based on detailed data on land use provided by Knol et al. (2004).[1]

Floor prices. We estimate floor prices conditional on housing characteristics from the universe of housing transaction data by the Dutch Association of Real Estate Brokers (NVM). We run a hedonic regression similar to Teulings et al. (2018) and Gaigné et al. (2022), where we regress log price per square meter on dwelling characteristics, neighbourhood, year and building-year dummies. The `NVM` data (years 2012 - 2021, $N = 1,324,690$), covering the large majority of owner-occupied housing transactions, include transaction prices, lot sizes, interior floor space size (m^2), addresses and housing attributes such as house type, number of rooms, construction year, garden, state of maintenance, central heating, and listed building status.

[1] For the 1900 land use maps, Knol et al. (2004) scanned and digitised historic maps into 50×50 metre grids, classifying them into 10 categories, such as built areas, water, sand, and forest. We aggregate these categories into three broad groups: built-up areas, open space, and water bodies.

A.3 Mode choice (OViN survey)

We use the 2016-2022 OViN survey for information on mode choice per neighbourhood combination. Within the OViN survey we consider only work commuting trips, *i.e.*, recreational trips or shopping trips are excluded. When trips are recorded as multi-modal trips, then we assign the trip to the mode that was used for the majority of the km traveled. Public transport covers trips made by train, bus, tram, metro and ferry. Trips by motorcycle or moped are discarded from the data. The number of person-commuting day observations in total is $N = 48,323$. The data contains trip origin (residence) and destination (workplace) neighbourhood, the modes used, the number of transfers, the date of surveyed respondents, sex, age, level of education and income, and self-reported travel distance and time.

A.4 Travel times

Travel times by road. We calculate travel times between neighbourhood pairs for different travel modes. Travel times are determined along the road, cycling, and pedestrian networks using the *Open Street Map* (OSM) database combined with assumptions on average travel speeds per mode and road type. We include all Dutch roads in the routing procedure, not just the roads for through traffic. The OSM database differentiates between road types in road infrastructure (motorways, trunks, primary, secondary, tertiary roads and residential roads) as well as the type of cycling infrastructure (on-road cycling lanes and separate cycleways).

Average speeds per mode and road type are set within the dodgr many-to-many routing algorithm. For example, the average speed of foot travel (on any walkable road) is assumed to be 5km/h. The speed of travel by bicycle (on any road) is on average 15km/h. Average car speeds depend on road type. On residential roads the average *free-flow* speed is 48km/h, on tertiary roads 60km/h, on secondary and primary roads 80km/h and motorways 112km/h. To route from (or to) a neighbourhood, we determine for each neighbourhood the intersection closest to its centroid, and then route from there.

In all three modes of travel on the road, the routing procedure yields free flow travel times in minutes for $12{,}782^2 \times 3 = 490{,}138{,}572$ origin-destination pairs. To reduce the data burden of our effort, we impute the travel time for long trips using the (calculated) mode-specific median travel speed along a straight line. We do this for trips for which it is unlikely that they will experience nonzero commuter flows. For car commuting, we do the imputation for trips longer than two hours, and for cycling and walking we use a one-hour threshold.

Congested car travel times. To perform our estimation of the congestion parameters in section 4.3, we require knowledge of the traffic delays of cars on both motorways and urban roads. To gauge these delays, we sample 1,000 neighbourhood pairs that are less than 100km apart and request congested driving times at 08:00 am (peak hour) on a regular weekday. We explain how we then predict congested travel times for all location pairs.

Travel times by public transport. We source information on public transport routes, schedules, and geographic public transport details from *General Transit Feed Specification* (`GTFS`) Schedule.[2] `GTFS` is a standardised open source data format that provides all the necessary ingredients to calculate the travel times of public transport between all combinations of neighbourhoods in our data.

The data for the Dutch public transport network can be accessed via a `GTFS` respository and API requests.[3] Using the `GTFS` timetables, we first determine the fastest routes between all Dutch public transport stops. Routes can include transfers and switching modes, when necessary. Then, we determine the fastest public transport routes between all pairs of neighbourhoods, where a commuter could in principle use any stop within a radius of 3 km. We assume that commuters travel at an average speed of 7.5 km/h along the road network when moving from a neighbourhood centroid to their intended public transport stop.[4]

A.5 Imputation methodology

There are two data sources for which we observe only a subset of neighbourhood combinations. These are mode shares (from `OViN`) and congested car travel times (from `Google`). We need mode shares for the entire neighbourhood domain to estimate congestion parameters $\{\lambda_1, \lambda_2\}$. We also need congested travel times for estimation of almost all relevant model parameters, $\{\nu, \kappa_1, \kappa, \lambda_1, \lambda_2\, \delta, \rho\}$. Hereafter, we explain our imputation methodology. The procedure is virtually identical for the imputation of mode shares and congested travel times. The former uses the 2016-2022 `OViN` mode shares as the training sample, while the latter uses congested travel times between 1000 neighbourhood combinations requested through the `Google API` as the training sample.

Mode shares Based on the `OViN` sample, we train a random forest model (see the `R` package `RandomForest`) on observed individuals in the survey using a 50% sample of the available

[2] General information on the data format can be found at https://gtfs.org/.

[3] From the repository website https://gtfs.ovapi.nl/nl/ we use the data of 2021, which is the earliest available year, closest to our base year 2016.

[4] An average speed of 7.5 km/h along the road network is in between common walking and cycling speeds. Dutch commuters often bicycle to train stations but commonly walk to bus and tram stops.

municipalities of residence. Next to mode-specific travel time, we add the following determinants of mode choice X_{ij} for both the residence and the work neighbourhood: travel time (per mode), distance from the nearest motorway ramp, distance to the nearest public transport stop, density and neighbourhood population. Finally, we add Euclidean distance between residence and work location ($dist_{ij}$). The choice of determinants is motivated by the fact that we observe them for all neighbourhood pairs, which guarantees that we can use the model to impute all necessary mode shares. Since the public transport and walking modes have much less observations than the cycling and car ones, we give the former higher weights in the training procedure. We estimate:

$$\Pr(\text{mode} = k|ij) = f\left(\tau_{ijm}, X_{ij}, dist_{ij}\right) \tag{A1}$$

We evaluate if the model is able to accurately predict the mode shares from any given municipality of residence to all associated work locations. To achieve this, we predict mode shares per residence-workplace-mode ijm on neighbourhood level. Then, we aggregate by residence population to construct mode shares per residence-mode on municipality level. The correlation between predicted and actual modes shares equal 0.97.

Congested car travel times Based on the `Google` sample, we train a random forest model (see the`R` package `RandomForest`) on observed neighbourhood combinations. As for the mode shares, we add determinants X for both residence and work neighbourhood, which include car free-flow travel times, distance to nearest motorway ramp, distance to nearest public transport stop, density and neighbourhood population. Furthermore, we add the Euclidean distance between residence and work location ($dist_{ij}$). All of these variables are observed for every neighbourhood combination, making them ideal inputs for the random forest model. Eventually, the model is used to generate predictions of congested travel times for all missing neighbourhood combinations, which can be formalised as follows:

$$\widehat{\tau_{ij1}} = f\left(\tau_{ijm}, X, dist_{ij}\right) \tag{A2}$$

The correlation between predicted and actual congestion delays equal 0.99.

B Structural estimation: Additional information and results

B.1 Calibrating aggregate mode shares

As substitution between modes plays a central role in our research question, we aim to match the aggregate mode shares implied by the model with the ones we observe in the `OViN` survey from the years 2016-2022 (compare Table 1, Panel A). Given the observed commuting flows N_{ij}, our Nested Logit mode choice is capable of reproducing the observed aggregate mode shares with precision.[1]

Although their distributions are similar, the commuting flows between location pairs implied by the model baseline differ from observed ones. The main reason for this is that the ASRW model cannot account for factors that lead to variability in mode choice *between* location pairs ij, apart from the generalised commuting time t_{ij}. Apart from t_{ij}, the probabilities of commuting depend on the attractiveness of the workplaces (transformed wages ω_j) and the residences (transformed amenities a_i), but not on their combinations.

To match observed aggregate mode shares, we recalibrate the alternative specific preference shifters b_m using a procedure outlined by Train (2009). First, we simulate the baseline configuration of the model for a single iteration and calculate the model implied mode shares $\hat{S}_m$.[2] We then update b_m in the following way, iterating until the sum of squared errors between b_m^1 and b_m^0 is sufficiently small:

$$b_m^1 = b_m^0 + \log \frac{S_m}{\hat{S}_m^0} \tag{B1}$$

Finally, we used the final recalibrated values b_m to run the baseline model until convergence. This ensures that the model matches, in the aggregate, the observed mode choices of the universe of Dutch workers.

B.2 Mode choice estimation

This subsection briefly discusses the instrumentation choice and first-stage results from estimating (29). Table B1 is separated into two panels. While Panel A refers to the first-stage results, Panel B repeats the second-stage results as presented in the main paper for comparison. In the single κ case (column 3) and the mode-specific κ_m case (column 6), we run the same first stage estimation and instrument mode-specific commuting time τ_{ijm} with Euclidean distance and its squared counterpart as is standard in the economic literature. As expected, both instruments are

[1] This is a property of Nested Logit estimation, see Train (2009), page 37 for a discussion on this matter.

[2] Relying on a single iteration saves computing time and does not yield materially different recalibration results compared to letting the model run until convergence in this step.

highly relevant, with a larger parameter magnitude for short-distance commuting modes (*i.e.*, walk). This is because the error of margin between Euclidean distance and calculated network travel times is smaller for shorter trips.

B.3 Congestion elasticity estimation

This subsection discusses briefly the instrumentation choice and first-stage results from estimating (29).

Recall that the congestion estimation uses information on car commuters to each workplace location. As the 2016-2022 `OViN` survey data is a representative sample of the entire Dutch population, we cannot directly infer the share of car commuting to each location of work from the data. Therefore, we obtain the share of car commuters by imputation as explained in Appendix A.5.

Table B2 shows the results from estimating (32). While Panel A refers to the first stage results, Panel B repeats the second stage results as presented in the main paper for comparison (compare (33) for a separation of the two stages). We employ a classical strategy used in the literature on agglomeration economies and use an historic instrument to address endogeneity concerns with car density. Decisions regarding workplace locations are likely influenced by the placement of infrastructure. However, we expect that for location decisions around 1900 this is not the case, as they are likely not correlated to current infrastructure supply, for example, because highways and roads suitable for cars did not exist around that time and commuting to work over longer distances was still uncommon.

Column (2) of Panel A regresses car density on the spatially weighted historic population density of 1900. Column (4) regresses car density on spatially weighted historic population density of 1900 and its squared counterpart. In both cases the instruments are highly relevant, with an F-statistic above 15,000. Controlling for the squared term in column (4), reveals a concavity in the data, as the parameter on the squared spatially weighted historic population density becomes negative. This implies that car density (and therefore congestion) is not necessarily the highest where historic city centres used to be.

B.4 Productivity and amenity estimation

This subsection discusses the instrumentation choice and first-stage results from estimating productivity and amenity parameters.

Table B3 shows the results from estimating (37) and (38). While Panel A refers to the first-stage results, Panel B repeats the second-stage results as presented in the main paper

Table B1 – Mode choice parameter estimation (incl. First Stage results)

Panel A: First Stage						
Dependent variable:	*Mode-specific travel time*					
	Single $\kappa_m = \kappa$			Mode-specific κ_m		
	(1)	(2)	(3)	(4)	(5)	(6)
Euclidean distance × Car			0.001***			0.001***
			(0.000)			(0.000)
Euclidean distance × Bicycle			0.005***			0.005***
			(0.000)			(0.000)
Euclidean distance × Walk			0.016***			0.016***
			(0.000)			(0.000)
Euclidean distance × PT			0.003***			0.003***
			(0.000)			(0.000)
(Euclidean distance)^2 × Car			0.000***			0.000***
			(0.000)			(0.000)
(Euclidean distance)^2 × Bicycle			0.000***			0.000***
			(0.000)			(0.000)
(Euclidean distance)^2 × Walk			0.000***			0.000***
			(0.000)			(0.000)
(Euclidean distance)^2 × PT			0.000***			0.000***
			(0.000)			(0.000)
Number of observations			130,240			130,240
First stage F-statistic			4,940			4,940
Panel B: Second Stage						
Dependent variable:	*Mode choice*					
	NL	NL	NL-CF	NL	NL	NL-CF
Commuting time elasticity, κ	-0.077***	-0.077***	-0.080***			
	(0.001)	(0.001)	(0.001)			
Car commuting time elasticity, κ_1				-0.097***	-0.094***	-0.070***
				(0.002)	(0.002)	(0.004)
Bicycle commuting time elasticity, κ_2				-0.093***	-0.093***	-0.089***
				(0.001)	(0.001)	(0.001)
Walk commuting time elasticity, κ_3				-0.128***	-0.128***	-0.148***
				(0.005)	(0.005)	(0.006)
PT commuting time elasticity, κ_4				-0.065***	-0.063***	-0.039***
				(0.002)	(0.002)	(0.003)
Non-car nest parameter, ψ_2	0.590***	0.603***	0.531***	0.849***	0.846***	0.972***
	(0.013)	(0.014)	(0.014)	(0.021)	(0.022)	(0.031)
Worker characteristics		✓	✓		✓	✓
First stage errors			✓			✓
Number of observations	130,240	130,240	130,240	130,240	130,240	130,240
Pseudo-R^2	0.293	0.295	0.304	0.307	0.317	0.320

Notes: PT stands for public transport. Data is sourced from the `OViN` Survey for the years 20216-2022 and consists of workers facing the choice set of four modes of transport (walk, bicycle, car, public transport). Worker characteristics include age, sex, education, income and day of the week. All columns estimate Nested Logit (NL) models. Car is the reference mode, meaning that $b_1 = 0$. As car is the only mode within the car nest it follows that $\psi_1 = 1$. Columns (1) and (2) estimate a baseline model with a single commuting time coefficient, implying one κ for all modes. Columns (3) and (4) use mode-specific commuting time coefficients κ_m. In columns (2) and (4), commuting time is instrumented with Euclidean distance for each mode using a control function approach (CF). Bootstrapped standard errors (100 replications) clustered at the i and j level in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

TABLE B2 – CONGESTION ELASTICITY
(INCL. FIRST-STAGE RESULTS)

Panel A: First Stage				
Dependent variable:	*Car density*			
	(1)	(2)	(3)	(4)
Spatially weighted population density in 1900		0.774***		1.089***
		(0.014)		(0.016)
(Spatially weighted population density in 1900)^2				-0.192***
				(0.007)
Panel B: Second Stage				
Dependent variable:	*Recovered workplace fixed effect,* $\log \hat{T}_{Mj}$			
	(1) OLS	(2) CF	(3) OLS	(4) CF
Car density, λ_1	0.013***	0.018***	0.013***	0.014***
	(0.000)	(0.001)	(0.000)	(0.000)
(Car density)^2, λ_2			0.004***	0.004***
			(0.000)	(0.000)
First stage error		✓		✓
Number of observations	12,782	12,782	12,782	12,782
First stage F-statistic		15,214		17,732

Notes: CF stands for 'control function'. Column (1) estimates log workplace fixed effects on car density following (32). In column (2), we instrument car density using historic population density in 1900. Column (3) estimates log workplace fixed effects on car density and car density squared. In column (4), we instrument car density with historic population density in 1900 and its squared counterpart. Bootstrapped standard errors (100 replications) clustered at the i and j level in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

for comparison. We instrument spatially weighted employment density and spatially weighted residential density, respectively, with spatially weighted historic population density in 1900 (see columns (2), (4) and (6)). Omitted variable bias concerns arise from the fact that the most attractive locations in terms of productivity and amenity levels may be correlated with unobserved location endowments. Unobserved endowments are likely uncorrelated over a very long time period, given that the Dutch economic geography has substantially changed over the last century. This makes historic population information arguably a suitable instrument, particularly given that we control for travel-to-work-area fixed effects in all specifications.

The results in Table B3 show that the elasticity to employment density and residential population density is negative, respectively. This implies that current jobs and residence locations are on average located outside of historic city centers. The spatial decay of historic population density is positive and larger when regressing on employment density. This indicates that the population density in 1900 is more locally concentrated than the current job density. The results are comparable if we assume amenities to diffuse via the car road network (column (4)) or the cycling network (column (6)). The difference in the magnitude of the spatial decay parameter is likely due to differences in speed between the two modes.

Table B3 – Productivity and amenity elasticities (incl. first stage results)

Panel A: First Stage						
Dependent variable:	*Spatially weighted employment density (log)*		*Spatially weighted residential density (log)*			
	Car travel time		Car travel time		Bicycle travel time	
	(1)	(2)	(3)	(4)	(5)	(6)
Spatially weighted population density in 1900 *(log)*, elasticity		-0.588***		-0.282***		-0.260***
		(0.001)		(0.022)		(0.022)
Spatially weighted population density in 1900 *(log)*, decay		0.092***		0.056***		0.018***
		(0.011)		(0.002)		(0.001)
Panel B: Second Stage						
Dependent variable:	*Productivities,* $\log \hat{A}_j$		*Amenities,* $\log \hat{B}_i$			
	(1) NLS	(2) NLS-CF	(3) NLS	(4) NLS-CF	(5) NLS	(6) NLS-CF
Agglomeration elasticity, γ	0.027***	0.023***				
	(0.001)	(0.001)				
Agglomeration decay, δ	-0.009***	0.078***				
	(0.001)	(0.001)				
Amenity elasticity, υ			0.037***	0.045***	0.078***	0.045***
			(0.004)	(0.006)	(0.003)	(0.006)
Amenity decay, ρ			0.007**	0.095***	0.378***	0.035***
			(0.002)	(0.010)	(0.024)	(0.003)
TTWA fixed effects	✓	✓	✓	✓	✓	✓
First stage error		✓		✓		✓
Number of observations	11,901	11,901	12,348	12,348	12,348	12,348
Pseudo-R^2	0.82		0.15		0.31	
First stage F-statistic		5,672		296		358

Notes: NLS stands for non-linear least squares, CF stands for control function, and TTWA stands for a travel-to-work-area. Column (1) estimates equation (37) by regressing local productivities on the spatially weighted employment density around j. In column (2), we instrument for the spatially weighted employment density using the population density in 1900. Column (3) estimates equation (38) by regressing local residential amenities on the spatially weighted residential density around i. In column (4), we instrument for the spatially weighted residential density using the population density in 1900. In columns (5) and (6), we repeat the estimation as in columns (4) and (5) assuming that amenity spillovers are governed by cycling travel times, as compared to car travel times. Bootstrapped standard errors (100 replications) clustered at the i and j level in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

C Counterfactual solution algorithm

To solve for the new equilibrium in each of the counterfactual scenarios, we largely follow the procedure described in Ahlfeldt et al. (2015) except for the computation of the mode choice probabilities and the addition of congestion as per Koster (2024). Initially, we select the starting values for transformed wages ω_j^0, amenities a_i^0, floor space prices p_i^0, and the total working population $\bar{H}$, corresponding to the values derived from the baseline scenario. We set travel times τ_{ijm}^1 for each mode to be in line with the counterfactual scenario. Subsequently, we compute the mode-specific commuting probabilities $\pi_{ij|m}$ using (15):

$$\pi_{m|k,ij}^1 = \sum_{n \in \mathbb{M}_k} \exp\left(\frac{b_n}{\psi_k} + \frac{\kappa_n}{\psi_k}\tau_{ijn}^1\right)$$

Now, we can update the mode nest inclusive value $t_{k|ij}^1$ using (14) and the generalized travel time t_{ij}^1 using (5):

$$t_{k|ij}^1 = \log \sum_{m \in \mathbb{M}_2} \exp\left(\frac{b_m}{\psi_k} + \frac{\kappa_m}{\psi_k}\tau_{ijm}^1\right)$$
$$t_{ij}^1 = \log\left[\sum_{k=1}^{2} \exp\left(\psi_k t_{k|ij}^1\right)\right],$$

and then, the generalized commuting probabilities π_{ij}^1 follow from (3) as:

$$\pi_{ij}^1 = \frac{a_i^0 \omega_j^0 \exp\left(-\eta\varepsilon t_{ij}^1\right)}{\sum_{r=1}^{I}\sum_{s=1}^{I} a_r^0 \omega_s^0 \exp\left(-\eta\varepsilon t_{rs}^1\right)}.$$

With the commuting probabilities we update the neighbourhood residential and workplace populations and their populations by mode:

$$H_{Ri}^1 = \sum_{s=1}^{I} \pi_{is}^1 \bar{H}$$
$$H_{Mj}^1 = \sum_{r=1}^{I} \pi_{rj}^1 \bar{H}$$
$$H_{Rim}^1 = \sum_{s=1}^{I} \pi_{ism}^1 \bar{H}$$
$$H_{Mjm}^1 = \sum_{r=1}^{I} \pi_{rjm}^1 \bar{H}.$$

In the next step, we update agglomeration and amenity externalities using (21) and (22):

$$A_j^1 = \bar{A}_j \left(\sum_{i=1}^{I} \frac{H_{Mi}^1}{(1-\iota_i)L_i} \exp(-\delta \tau_{ij1}^1) \right)^{\gamma}$$
$$B_i^1 = \bar{B}_i \left(\sum_{j=1}^{I} \frac{H_{Rj}^1}{(1-\iota_j)L_j} \exp(-\rho \tau_{ij1}^1) \right)^{\upsilon}.$$

The share of commercial floor space θ_i^0 equals:

$$\theta_i^0 = F_{Mj}^0 / F_i^0$$

Using (24) optimal capital use K_1^0 becomes:

$$K_i^0 = \left(\mu \Upsilon_i p_i^0 \right)^{\frac{1}{1-\mu}} L_i^0$$

and floor space supply by the construction sector updates to

$$F_i^1 = \Upsilon_i K_i^{\mu} \left((1-\iota_i) L_i \right)^{1-\mu}$$
$$F_{Ri}^1 = (1-\theta_i^0) F_i^1$$
$$F_{Mi}^1 = \theta_i^0 F_i^1.$$

After updating workplace populations and floor spaces, we update production y_j^1 using (19):

$$y_j^1 = A_j^1 \left(\varsigma H_{Mj}^1 \right)^{\alpha} \left(F_{Mj}^1 \right)^{1-\alpha}$$

so that the wages w_j^1 become

$$w_j^1 = \frac{\alpha y_j^1}{\varsigma H_{Mj}^1}.$$

Then, we determine the expected wage at the residence $\mathbb{E}[\bar{w}_i^0]$ as

$$\mathbb{E}[\bar{w}_i^0] = \sum_{s=1}^{I} \pi_{is|i} w_s^0$$

and the floor space prices p_i^1 as

$$p_i^1 = \begin{cases} \dfrac{(1-\alpha)y_j^1}{F_{Mi}^1}, & \text{if } H_{Mj}^1 > 0 \\ \dfrac{(1-\beta)\mathbb{E}[\bar{w}_i^1]}{F_{Ri}^1}, & \text{if } H_{Mj}^1 = 0 \text{ and } H_{Ri}^1 > 0 \end{cases}$$

and the share of commercial floor space θ_i^1 updates to

$$\theta_i^1 = \frac{F_{Mi}^1}{F_{Ri}^1 + F_{Mi}^1}.$$

Now, we can use (4) to arrive at updated transformed wages ω_j^1 and transformed amenities a_i^1. In the last step, we update the congested car travel times τ_{ij1}^1 using (16):

$$\tau_{ij1}^1 = \tau_{ij1}^f T_{Ri} T_{Mj} \exp(\lambda_1 \Lambda_{Mj}^1 + \lambda_2 (\Lambda_{Mj}^1)^2),$$

where as in (17) we have the car traffic density

$$\Lambda_{Mj}^1 = \frac{\sum_{i=1}^{I} H_{Ms1}^1 \exp(-\kappa_1 \tau_{ij1}^0)}{\sum_{i=1}^{I} R_{s1} \exp(-\kappa_1 \tau_{ij1}^0)}.$$

Finally, we calculate the percentage changes per neighbourhood from the current iteration to the previous one. If these changes hit some threshold, the algorithm converges and we obtain a new model solution, otherwise we continue iterating. Usually, this happens in about 12-15 iterations or 3-6 hours on a normal machine. As in Ahlfeldt et al. (2015), we update the endogenous variables $\{\omega_j, a_i, p_i, \tau_{ij1}\}$ using an update factor. After some experimentation, a factor of 0.7 or 70% weight on the previous iteration value (and 30% on the current iteration) appears to work best for convergence.

D Counterfactual analysis: Additional information and results

D.1 Removing separate cycleways

In the *No separate cycleways* counterfactuals, we remove cycling infrastructure elements that were present in the reference situation if they satisfy a number of criteria based on road types, maximum speeds, traffic signs and number of lanes as indicated in `OSM`. With these criteria, we aim to only remove separate cycleways situated next to roads. Table D1 gives a detailed overview of these criteria.

In `OSM`, one can model cycling infrastructure in three ways.[1] First, they can use a dedicated road type `cycleway`. Alternatively, they can give the value to the `bicycle` tag of a road to indicate that the road has a cycling track next to it that is clearly separated from the main road. If cycling is not permitted on a road altogether, the `bicycle` tag is set to `no`. Finally, they can indicate on-road cycling provisions such as cycling lanes and cycling streets using tags like `cycleway=lane`.[2] In our effort, we will focus solely on separate cycleways and leave on-road cycling provisions in place.

Our approach proceeds in three steps. We first remove the `OSM cycleway` road type in its entirety. Subsequently, we disallow on-road cycling on primary and secondary roads that feature the `use_sidepath` flag or that have 2 or more lanes in a single direction. Then, we do the same for roads where the maximum speed is greater than 60 km/h or where the existing cycleway is compulsory from a legal perspective.[3] In all other cases, we leave the cycling infrastructure unchanged with respect to the reference situation.

After having removed the separate cycleways, we reassign the these lanes to cars in particularly dense areas (*i.e.*, neighbourhoods with more than 1000 addresses per km^2). This affects our measure of car road lane kilometer per neighbourhood (R_j), which we recalculate for the *No cycling* and *No separate cycleways* scenarios, respectively. Thereby we count one separate cycleway to the right and to the left of the road as one additional car lane each.

D.2 Counterfactual results: Additional maps

In the main text, we discuss percentage changes in residential density (for all counterfactuals) and residential floor space consumption (for the *No cycling* scenario) using maps at the neighbourhood

[1] See `https://wiki.openstreetmap.org/wiki/Bicycle` for a complete specification and overview of `OSM` bicycle infrastructure tagging rules and guidelines.

[2] `OSM` features a fourth cycling infrastructure tag type `track`, but it is not used in the Netherlands (see `https://taginfo.geofabrik.de/europe/netherlands/keys`).

[3] In `OSM`, this is indicated by traffic sign codes `NL:G11` and `NL:G12a`. As the traffic sign `NL:G1` indicates a non-compulsory cycleway, we assume cyclists will cycle on the road in that case.

Table D1 – Criteria for removal of cycling infrastructure in the *No separate cycleways* scenario

On-road cycling is not allowed ...	Relevant OSM tags
... if cycleway use was compulsory	`traffic_sign=NL:G11` `traffic_sign=NL:G12a`
... on primary and secondary roads with more than one lane in either direction	`lanes >= 2` `bicycle == "use_sidepath"`
... on any road with a maximum speed >60 km/h	`max_speed > 60`
...on separate cycleways themselves. Cycleway road type is completely removed. Cyclists cycle on road only if the other rules permit it.	`highway=cycleway`

level. In this section, we provide additional maps for workplace density, wages net of commuting and floor prices for all scenarios.

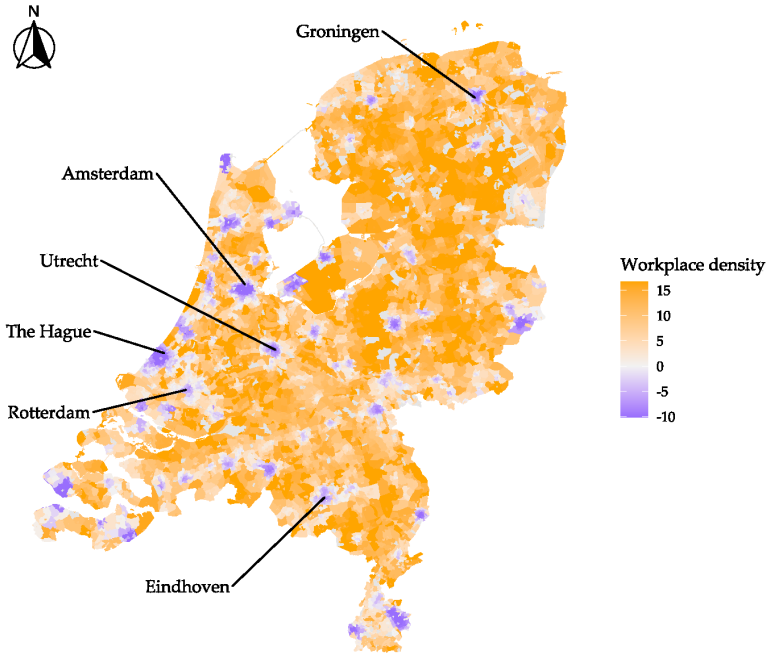

(a) Workplace density Δ%

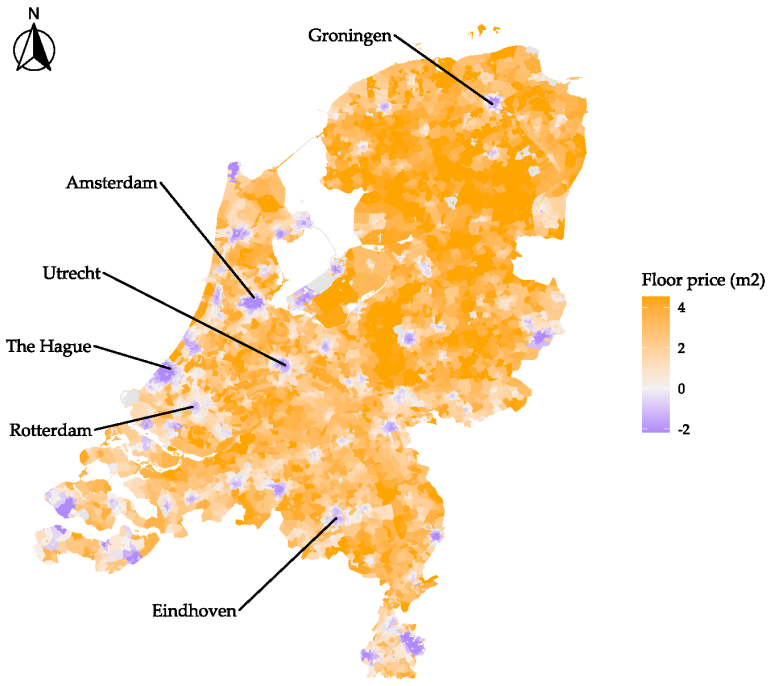

(b) Floor prices Δ%

Notes: This figure illustrates the local spatial differences arising from analysing the *No cycling* scenario. The upper figure (a) shows changes in workplace density (in percent), while the lower figure (b) shows changes of floor prices (in percent). For contrast, percentages displayed are held between the 5^{th} and 95^{th} percentile. Both figures display the names and locations of the six largest cities in the Netherlands.

Figure D1 – *No cycling* scenario

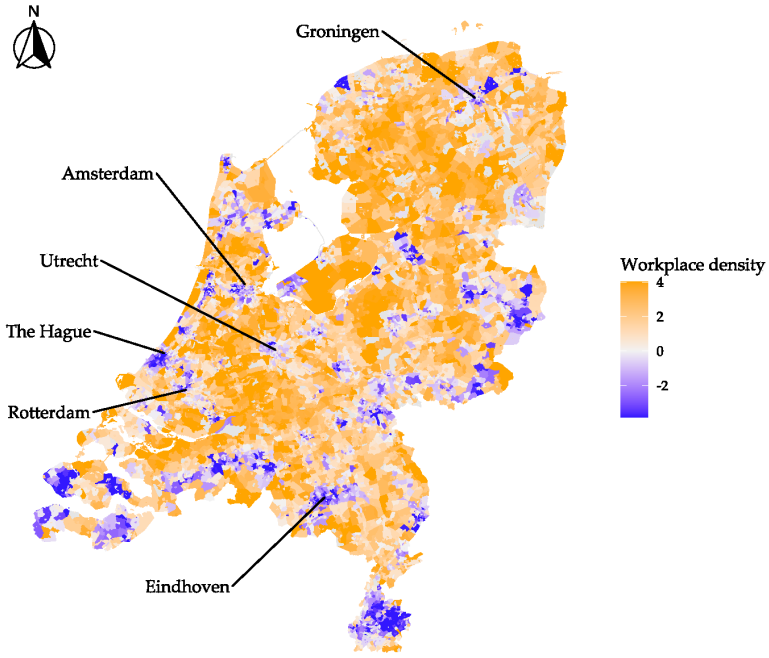

(a) Workplace density Δ%

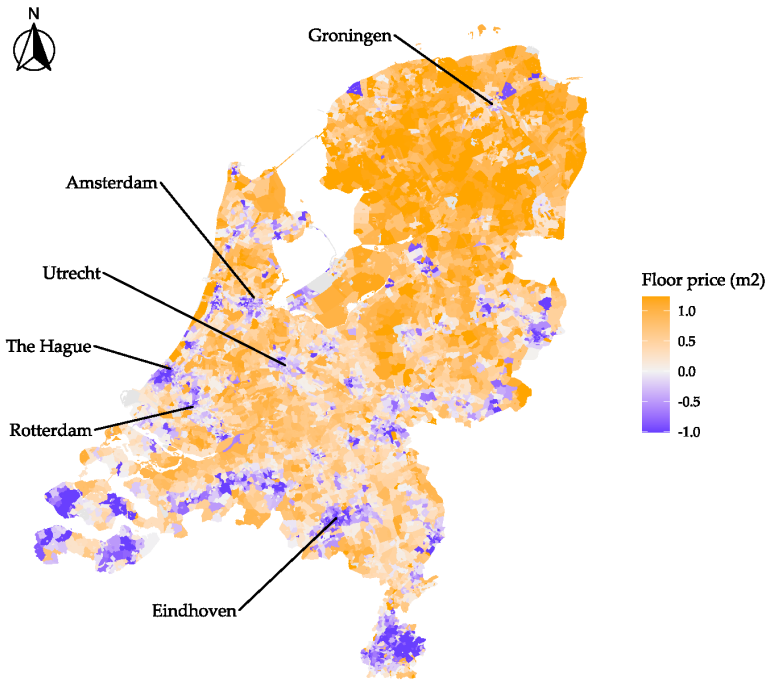

(b) Floor prices Δ%

Notes: This figure illustrates the local spatial differences arising from analysing the *No cycleways* scenario. The upper figure (a) shows changes in workplace density (in percent), while the lower figure (b) shows changes of floor prices (in percent). For contrast, percentages displayed are held between the 5^{th} and 95^{th} percentile. Both figures display the names and locations of the six largest cities in the Netherlands.

Figure D2 – *No cycleways* scenario

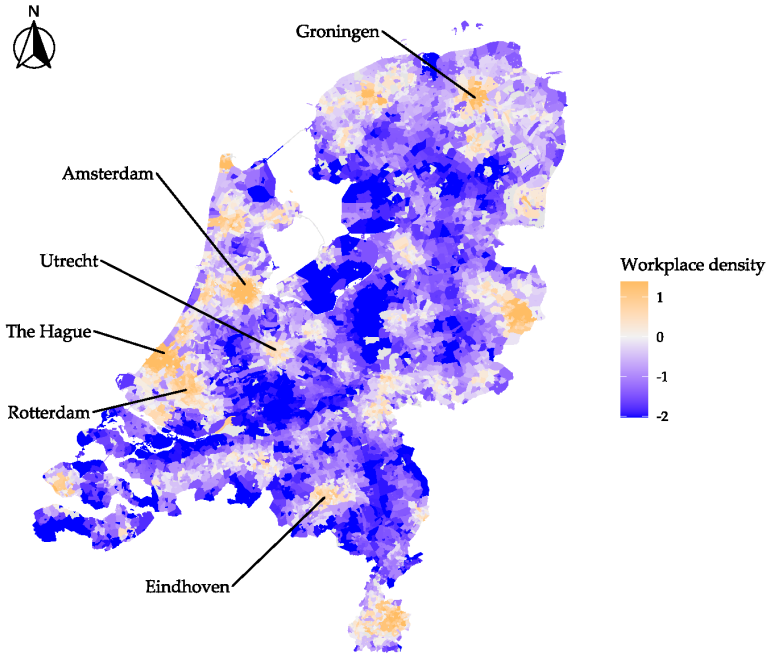

(a) Workplace density Δ%

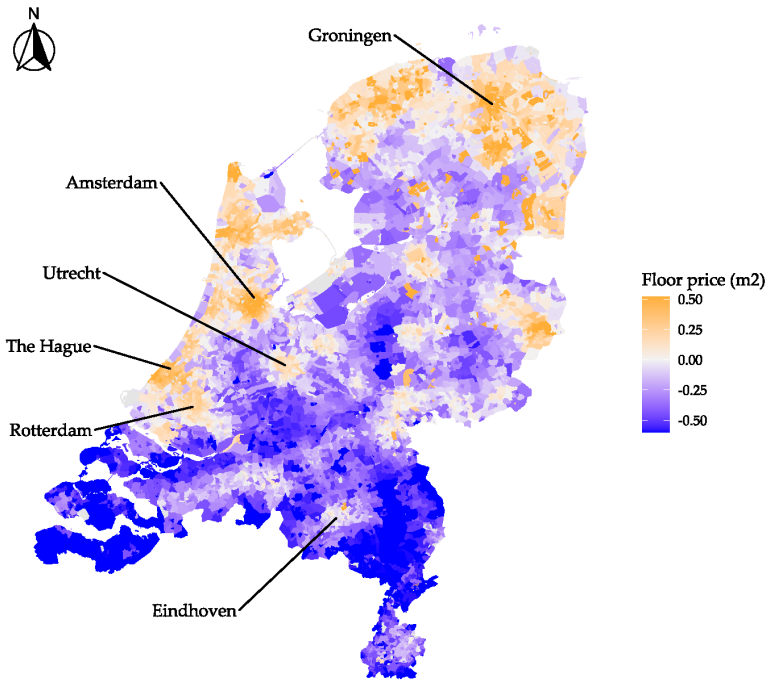

(b) Floor prices Δ%

Notes: This figure illustrates the local spatial differences arising from analysing the *Cycling speed increase* scenario. The upper figure (a) shows changes in workplace density (in percent), while the lower figure (b) shows changes of floor prices (in percent). For contrast, percentages displayed are held between the 5^{th} and 95^{th} percentile. Both figures display the names and locations of the six largest cities in the Netherlands.

Figure D3 – *Cycling speed increase* scenario