Beyond distance: modelling cyclists' route choice using calibrated empirical travel time

Laurent Cazor*[1], Corin Edward Staves[1], Mads Paulsen[1], and Thomas Kjær Rasmussen[1]

[1]Department of Technology, Management and Economics, Technical University of Denmark, Akademivej Bygning 358, 2800 Kgs. Lyngby, Denmark

Short summary

This paper develops a cyclist route choice modelling framework based on calibrated empirical travel time rather than distance alone or conventional infrastructure proxies. Using a large crowdsourced GPS dataset from Copenhagen, Denmark, we identify cruising segments in smoothed and map-matched trajectories, estimate non-parametric empirical link-level cruising speed distributions, and regularise them through hierarchical Bayesian shrinkage. We then estimate link-level delays as the difference between observed traversal time and cyclist-specific expected cruising time, again using shrinkage to ensure network-wide coverage. These cyclist- and link-specific travel times are aggregated to the route level and introduced into multinomial logit route choice models. The results show that travel-time-based specifications substantially outperform conventional distance-based models, and that the best-performing linear specification distinguishes between cruising time and delay. The framework provides a basis for future extensions toward more fundamental route costs, including, e.g., physical effort.

Keywords: bicycle route choice, calibrated travel time, GPS trajectories, Bayesian shrinkage, delays

1 Introduction

Cyclists' route choice is often modelled using route distance together with infrastructure attributes such as facility type, intersection density, road surface, and slope (Łukawska, 2024). While these variables explain much of the variation in observed choices (e.g. Łukawska et al., 2023; Meister et al., 2023; Broach et al., 2012), they typically act as proxies for underlying behavioural factors. For instance, signalized intersections generate time losses through deceleration, stopping, waiting, and re-acceleration, whereas protected facilities may be valued partly because they allow higher and more stable speeds (Bigazzi & Lindsey, 2019). This suggests that cyclists' choices respond not only to infrastructure categories but also to the travel-time implications of the riding environment.

Models relying heavily on dummy-coded infrastructure attributes can be difficult to interpret and may transfer poorly across networks with different classification systems. Recent work therefore emphasises more fundamental representations of route costs (Łukawska, 2024), in line with route choice models for motorised travel where travel time plays a central role.

One notable effort to estimate bicycle route choice in value-of-time (VoT) space is the Gothenburg study by Liu et al. (2025). Their model used a trip-level desired speed and intersection-type average delays, providing a useful proof of concept but relying on coarse assumptions. As a result, the VoT specification performed markedly worse than the corresponding value-of-distance (VoD) model, indicating that overly aggregated travel-time representations can degrade behavioural realism. This underlines the need for a more refined treatment of cyclists' speed heterogeneity and local operating conditions.

In this paper, we develop a route choice framework in which cyclists evaluate routes based on

*Corresponding author: lauca@dtu.dk

calibrated link travel times derived from large-scale crowdsourced GPS trajectories. We identify cruising segments, estimate non-parametric cruising-speed distributions for each network link, and regularise sparsely observed links through hierarchical Bayesian shrinkage across infrastructure, environment, surface, and slope classes. Delays are then estimated by comparing observed traversal times with cyclist-specific expected cruising times and shrunk analogously.

The resulting framework yields cyclist- and link-specific travel times across the entire network. This enables a comparison of distance-based models, time-based models, and a decomposed specification distinguishing cruising time and delay. The paper contributes by replacing infrastructure proxies with costs tied directly to observed riding conditions and by demonstrating that detailed, individual-specific travel-time components improve both explanatory power and behavioural interpretability, offering a foundation for future extensions including energetic costs.

2 Data and preprocessing

The empirical analysis uses a large crowdsourced GPS dataset of bicycle trips in Copenhagen and Frederiksberg (Łukawska et al., 2023). The data contain 234,860 trajectories from 5,451 cyclists, covering a wide range of cycling environments, facility types, and street categories.

Raw trajectories were first smoothed and temporally interpolated to an approximate resolution of one point per second, then map-matched to a detailed OpenStreetMap-based cycling network using a Hidden Markov Model (Newson & Krumm, 2009). These preprocessing steps are taken as given; the present work uses the resulting smoothed and map-matched trajectories to estimate cruising speeds and link-level delays.

The matched network was enriched with link attributes derived from OpenStreetMap and elevation data, including infrastructure type, road environment, surface quality, slope category, and link length. These attributes define hierarchical link classes used in the shrinkage procedure and support interpretation of the estimated travel-time components.

3 Method

The proposed framework constructs cyclist-specific route travel times by combining two components: an individual-specific expected cruising time on each link and an expected link delay term capturing stop-and-go conditions and other local frictions.

Because route choice models require travel-time estimates not only on observed links but also on the full set of feasible alternatives, the framework produces calibrated speeds and delays across the entire network. Hierarchical shrinkage ensures stable estimates even on sparsely observed links.

The procedure transforms raw GPS trajectories into link-level, cyclist-specific travel times through four main stages: (i) identification of cruising events from smoothed and map-matched trajectories, (ii) estimation of non-parametric link-level cruising-speed distributions with hierarchical shrinkage, (iii) derivation of cyclist-specific speed profiles from within-link relative performance, and (iv) calibration of link delays based on observed traversal times. Figure 1 summarizes the full pipeline and provides a visual overview of how these components interact. Each step is summarised in the following sub-sections.

Notation

Let A denote the set of links in the cycling network. For each link $a \in A$, let $c(a)$ denote its assigned class. Classes are defined by combinations of observable link attributes, including:

- **Road type:** Residential, Medium, or Large road; Cycleway; Pedestrian zone; Living street; etc.
- **Cycling infrastructure type (optional):** Separated track, Painted lane, or No dedicated infrastructure.
- **Surface type:** Smooth (asphalt), Gravel, or Cobblestones.

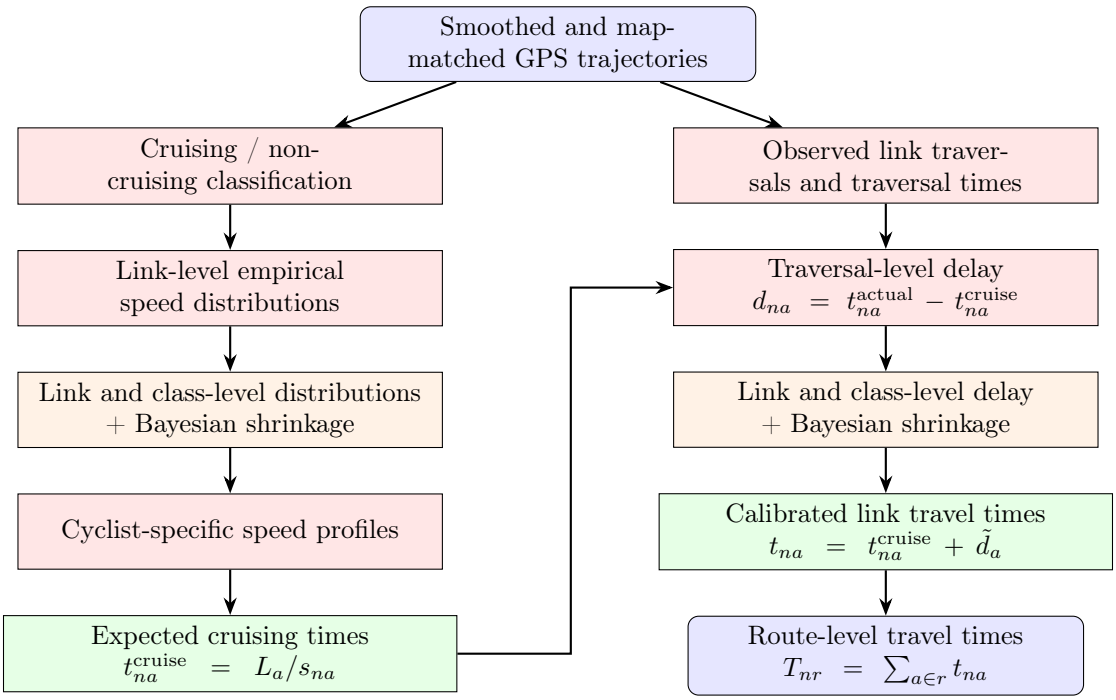


Fig. 1: Pipeline of the calibrated travel-time model.

- **Slope category:** Five groups ranging from very steep downhill to very steep uphill.

The Cartesian product of these attributes yields 240 possible class combinations. These classes are used both for hierarchical shrinkage of cruising-speed and delay estimates and for interpreting link-level heterogeneity in the calibrated travel-time model.

Cruising point identification

Raw bicycle trajectories include both cruising and non-cruising behaviour such as starting, stopping, and acceleration or braking. To recover cyclists' cruising speed while modelling delays separately, GPS points are classified into cruising and non-cruising phases. The classification follows Ma & Luo (2016) and uses local kinematic patterns (speed, acceleration, elapsed time, and travelled distance) to detect sustained stable-riding segments. Very low-speed points, initial acceleration phases, and marked acceleration or deceleration episodes are labelled non-cruising. All remaining points are treated as cruising. This preprocessing isolates trajectory segments that best reflect cyclists' cruising pace.

Link-level cruising-speed distributions

Cruising points are grouped by observed link traversals. For each traversal h of link a, we compute a traversal-level cruising speed $\bar{s}_{ha}$ as the mean of all cruising points on that link. The set of traversal speeds for link a is then represented through empirical quantiles

$$Q_a(q_p) = \text{Quantile}_{q_p}(\bar{s}_{1a}, \ldots, \bar{s}_{N_a a}),$$

where N_a is the number of traversals and $(q_1, \ldots, q_P)$ a fixed grid of quantile levels[1]. This non-parametric representation captures heterogeneity in link-specific cruising speeds without imposing a distributional assumption.

Hierarchical shrinkage across link classes

Many links have few (or no) traversal observations, making purely link-specific quantiles unstable. To ensure robust, network-wide coverage, links are grouped into hierarchical classes defined by infrastructure type, road environment, surface quality, and slope. Class-level empirical cruising-speed

[1] In the current study, we used 10 quantiles: (0.05, 0.15, ..., 0.95)

distributions are computed from pooled observations, and sparse links fall back to progressively coarser classes. Minimum data requirements are enforced throughout: at least 5 cyclists for cyclist–link profiles, 10 cruising observations for empirical link quantiles, and 200 observations for class-level distributions, with sparser sets handled through hierarchical shrinkage.

Let $Q_{c(a)}(q_p)$ be the quantile for the best-supported class of link a. Final link-level quantiles are obtained via shrinkage:

$$Q_a^{\text{final}}(q_p) = w_a^s Q_a(q_p) + (1 - w_a^s) Q_{c(a)}(q_p), \qquad w_a^s = \frac{N_a}{N_a + \kappa_s},$$

with $\kappa_s = 20$. Well-observed links retain their empirical distribution, with a class-specific weight near zero, while sparse links borrow speed information from similar links.

Cyclist-specific speed profiles

Although link-level quantiles describe plausible cruising speeds, cyclists differ systematically in their typical cruising pace. We therefore derive cyclist-specific speed profiles based on their relative performance across the links they traverse.

Cruising observations are first aggregated to the cyclist–link level. For cyclist n on link a, the length-weighted mean cruising speed is

$$\bar{s}_{na} = \frac{\sum_{m \in \mathcal{O}_{na}} s_{nam}\, \ell_{nam}}{\sum_{m \in \mathcal{O}_{na}} \ell_{nam}}.$$

Only links observed for at least a minimum number of distinct cyclists are retained. Cyclists are then ranked within each link to obtain percentile scores

$$q_{na} = \text{RankPct}_a(\bar{s}_{na}),$$

indicating how fast cyclist n rides relative to other cyclists on that link.

These percentiles are aggregated into a cyclist-level score

$$\phi_n = \frac{\sum_{a \in \mathcal{A}_n} \omega_{na} q_{na}}{\sum_{a \in \mathcal{A}_n} \omega_{na}},$$

where ω_{na} weights observations by traversed length and number of points. Cyclists are re-ranked using ϕ_n and assigned to quantile groups. Let θ_n denote cyclist n's quantile; expected cruising speed on link a is

$$s_{na} = Q_a^{\text{final}}(\theta_n), \qquad t_{na}^{\text{cruise}} = \frac{L_a}{s_{na}}.$$

This yields cyclist-specific cruising times consistent with both their riding pace and local link characteristics.

Link-level delay estimation

Observed traversal times are reconstructed from consecutive map-matched GPS points on the same link. Let t_{ha}^{start} and t_{ha}^{end} denote the first and last timestamps of traversal h on link a. To mitigate discretisation bias caused by GPS sampling, temporal gaps to the previous and next traversed links are split between neighbouring segments, and incomplete traversals are discarded. The resulting observed traversal time is

$$t_{ha}^{\text{actual}} = (t_{ha}^{\text{end}} - t_{ha}^{\text{start}}) + \frac{g_{ha}^{-}}{2} + \frac{g_{ha}^{+}}{2}.$$

Delay is defined as the difference between observed and expected cruising time:

$$d_{na} = t_{na}^{\text{actual}} - t_{na}^{\text{cruise}}.$$

This *reduced-form* quantity captures stopping, waiting, and other local slowdowns without identifying them separately.

Delays are averaged at the link level. Let $\bar{d}_a$ be the mean delay for link a, and $d_{c(a)}$ the corresponding class-level delay. Final delay estimates follow the same shrinkage structure as speeds:

$$\tilde{d}_a = w_a^d \bar{d}_a + (1 - w_a^d) d_{c(a)}, \qquad w_a^d = \frac{M_a}{M_a + \kappa_d},$$

where M_a is the number of delay observations and $\kappa_d = 20$. The calibrated travel time for cyclist n on link a is

$$t_{na} = t_{na}^{\text{cruise}} + \tilde{d}_a.$$

The delay term is modelled as link-specific rather than cyclist-specific, as stopping for signals, yielding, and queuing typically affect all cyclists on a link similarly. While individual riders may differ slightly in how often or how sharply they decelerate, these effects are modest relative to infrastructure-driven delay components.

Route-level travel time

For any route alternative r available to cyclist n, expected route travel time is obtained by summing calibrated link travel times:

$$T_{nr} = \sum_{a \in r} t_{na}. \tag{1}$$

These route-level travel times, including cruising and delay times, are then introduced into the route choice model as route cost attributes, allowing comparison with more conventional specifications based on distance and infrastructure-related proxies.

Route choice model

For each observed trip, a set of feasible alternatives $\mathcal{R}_n$ is generated using the combined labelling–stochastic method of Cazor et al. (n.d.), and route choice is estimated with a multinomial logit model. For cyclist n and route $r \in \mathcal{R}_n$, utility is

$$U_{nr} = \mu\left(C_{nr} + \sum_{k=1}^{K} \beta_k X_{nkr}\right) + \beta_{\text{PS}} \ln(\text{PS}_{nr}) + \varepsilon_{nr},$$

where C_{nr} is the main route-cost term, X_{nkr} the exposure to attribute k, PS_{nr} the path-size factor (Ben-Akiva & Bierlaire, 1999), and ε_{nr} an i.i.d. extreme-value term. The path-size factor is

$$\text{PS}_{nr} = \sum_{a \in r} \frac{t_{na}}{T_{nr}} \frac{1}{\sum_{j \in \mathcal{R}_n} \mathbf{1}_{aj}},$$

where $\mathbf{1}_{aj} = 1$ if link a belongs to route j and T_{nr} is expected travel time.

We compare several definitions of C_{nr} and X_{nkr}. In the baseline value-of-distance (VoD) specification,

$$C_{nr} = L_r = \sum_{a \in r} l_a, \qquad X_{nkr} = \sum_{a \in r} l_a \delta_{ak},$$

with l_a the link length and δ_{ak} an indicator for attribute k. In the value-of-time (VoT_{exp}) specification based on expected travel time,

$$C_{nr} = T_{nr} = \sum_{a \in r} t_{na}, \qquad X_{nkr} = \sum_{a \in r} t_{na} \delta_{ak},$$

with t_{na} depending on cyclist-specific cruising speed. A cruising-time specification replaces t_{na} by cyclist-specific cruising time:

$$C_{nr} = T_{nr}^{\text{cruise}} = \sum_{a \in r} t_{na}^{\text{cruise}}, \qquad X_{nkr} = \sum_{a \in r} t_{na}^{\text{cruise}} \delta_{ak}.$$

Finally, a decomposed specification separates cruising time and delay:

$$U_{nr} = \beta_{t_{no}} T_{nr}^{\text{cruise}} + \beta_d D_r + \sum_{k=1}^{K} \beta_k X_{kr} + \beta_{\text{PS}} \ln(\text{PS}_{nr}) + \varepsilon_{nr},$$

where

$$D_r = \sum_{a \in r} \tilde{d}_a, \qquad X_{kr} = \sum_{a \in r} l_a \delta_{ak}.$$

These formulations allow us to assess whether route choice is better explained by distance-based proxies, cyclist-specific time costs, or a decomposition into cruising and delay components.

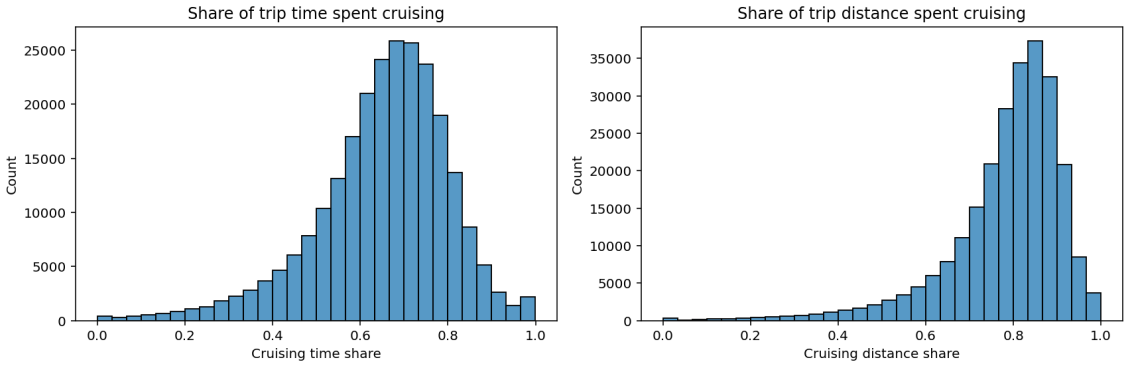
4 Results

Speed model

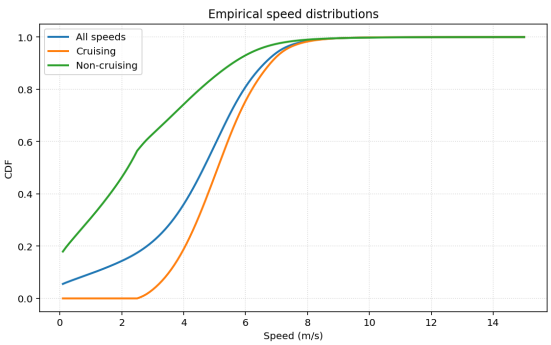
We first illustrate key outputs of the calibrated speed and travel-time model. Figure 2a reports the distribution of the share of cruising time and distance across trips, based on the cruising-event detection procedure. Cyclists spend on average 65% of their total travel time in cruising conditions, but over 78% of their travelled distance. This gap reflects the importance of interruptions, confirming that a substantial fraction of travel time is spent in stop-and-go or constrained conditions rather than cruising.

Figure 2b shows the empirical cumulative distribution functions (CDFs) of instantaneous speeds, distinguishing observations classified as cruising (69% of all GPS points) from those labelled as non-cruising. Cruising observations cluster at higher speeds with a lower variance, whereas non-cruising observations exhibit lower speeds associated with starting, stopping, and deceleration phases.

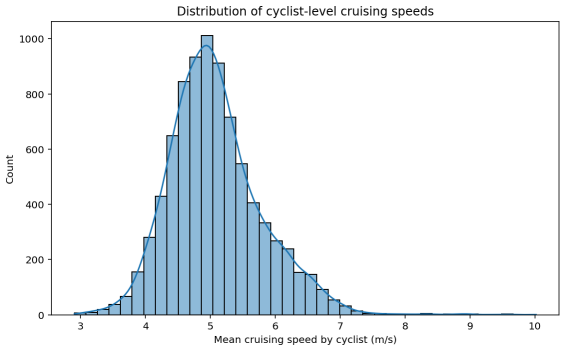
Figure 2c displays the distribution of individual average cruising speeds (computed over each cyclist's cruising observations). The mean cruising speed is approximately 5.11 m/s, but the distribution is wide, indicating substantial inter-individual heterogeneity. This motivates the cyclist-specific speed profiling used later in the travel-time model, as cyclists differ systematically in their preferred cruising pace.



(a) Distribution of cruising time and distance shares across trips



(b) CDF of cruising and non-cruising speeds



(c) Distribution of cyclists' mean cruising speeds

Fig. 2: Overview of cruising behaviour and speed distributions.

Figure 3 presents the resulting length-normalized delay estimates for central Copenhagen. Delays are highest around signalized intersections, where cyclists commonly stop and wait, and in pedestrian-heavy or mixed-use areas where cyclists must slow down or dismount. These spatial patterns illustrate the role of local frictions in shaping experienced travel times and highlight the need to represent delay explicitly within route choice models.

Figure 4 illustrates how the hierarchical shrinkage procedure combines link-specific and class-level cruising-speed distributions. For each example link, the figure shows three CDF curves: (i) the empirical link-specific distribution derived from its traversal-level cruising speeds, (ii) the corresponding class-level distribution, and (iii) the final shrunk distribution.

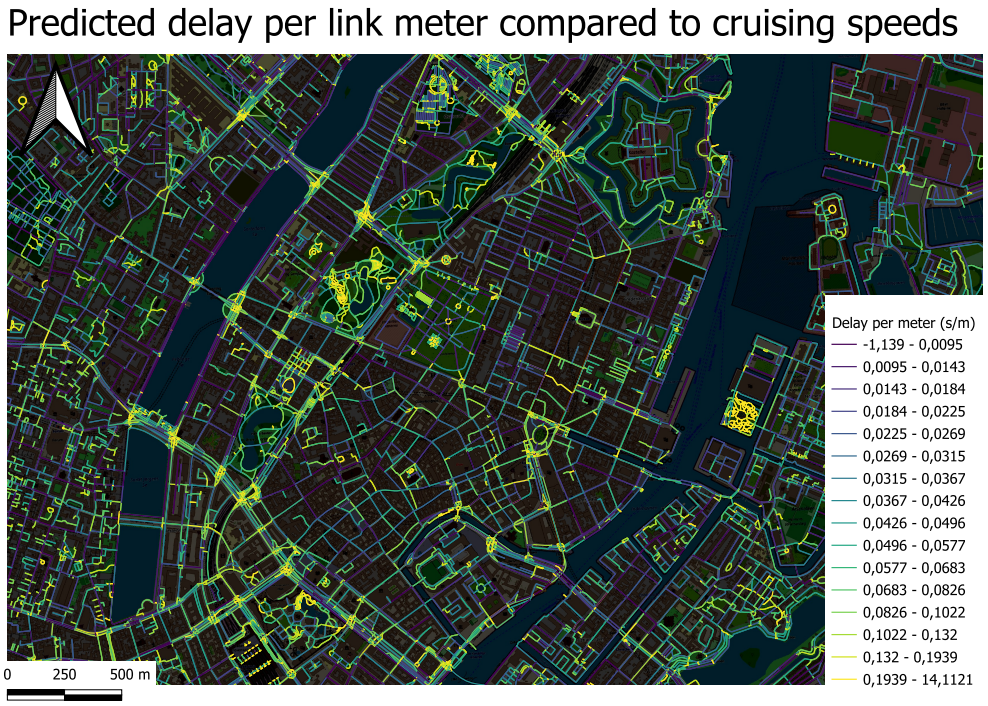
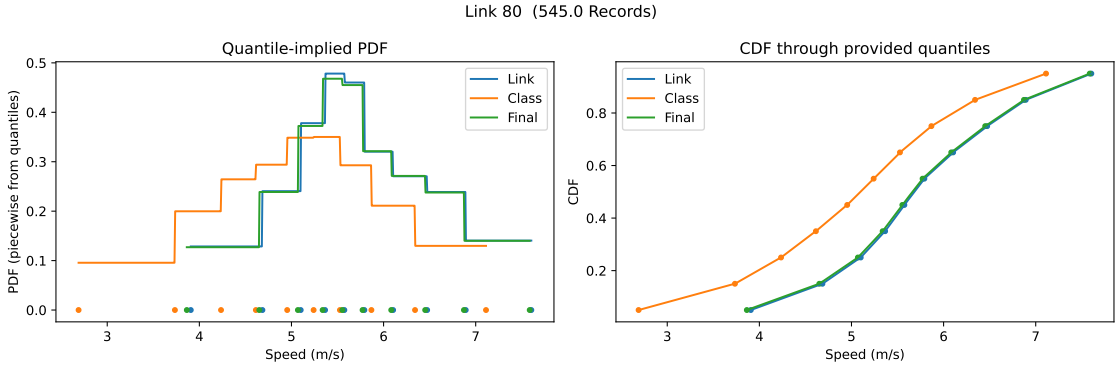
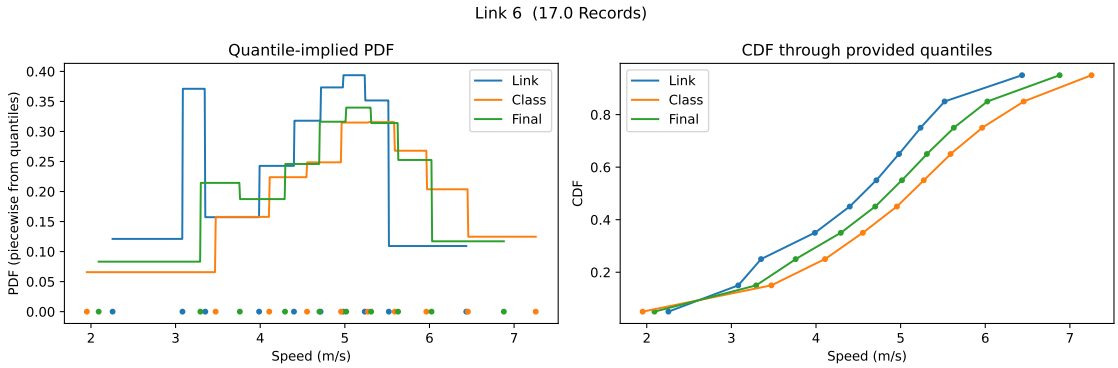


Fig. 3: Map of length-normalized delays on the network



(a) Link with 545 cruising observations ($w_a = 0.964$)



(b) Link with 17 cruising observations ($w_a = 0.459$)

Fig. 4: Example cruising-speed distributions for a well-observed link and a sparsely observed link. Each panel shows the empirical link-specific distribution, the class-level distribution, and the final shrunk distribution.

Figure 4a shows a link with 545 cruising observations, for which the shrinkage weight is very high ($w_a = 545/565 = 0.964$). Accordingly, the final distribution almost coincides with the link-

specific one. In contrast, Figure 4b shows a link with only 17 observations ($w_a = 17/37 = 0.459$). Here, the final distribution is visibly pulled toward its class-level counterpart, illustrating how shrinkage stabilizes estimates on sparsely observed links while preserving link-level information when sufficient data exist.

As a validation, a cruising-time-only model underestimated observed trip durations by roughly 20%, whereas the calibrated model combining cruising speeds and link-level delays reduced this bias to 2.9% on average. Remaining differences mainly reflect trip starts, ends, and occasional long pauses that are not intended to be captured at the link level.

Route choice model

Tab. 1: Model estimates (MNL). Ordered from worst to best fit (by Final LL). All estimates are significant at the 0.01 level unless noted by a †. Infrastructure and surface coefficients are normalized by the generic route-cost coefficient of each model: $\beta_k/\beta_{\text{length}}$ for VoD-type models, $\beta_k/\beta_{t_{no}}$ for VoT_{no}, and $\beta_k/\beta_{t_{exp}}$ for VoT_{exp}.

MNL specification **Parameter**	**Linear** **(VoD)**	**Linear** **(VoT_{no})**	**Linear** **(VoT_{exp})**	**Linear** **(VoD + t_{exp})**	**Linear** **(VoD + t_{no}, delay)**
Generic route-cost coefficients					
Length	-9.981	–	–	-3.370	-4.197
Cruising time (t_{no})	–	-2.387	–	–	-1.444
Expected time (t_{exp})	–	–	-2.083	-1.752	–
Delay	–	–	–	–	-2.003
ln(PS)	0.842	0.593	0.609	0.849	0.846
Normalized infrastructure and surface coefficients					
No infrastructure					
Residential Road	0.158	0.061	0.135	0.432	0.384
Medium Road	0.147	0.104	0.105	0.340	0.285
Large Road	0.352	0.419	0.388	0.969	0.783
Painted lane					
Residential Road	-0.084	-0.189	-0.077	-0.149	-0.081
Medium Road	0.034	-0.019	0.022	0.080	0.082
Large Road	-0.571	-0.666	-0.587	-1.406	-1.144
Separated Track					
Residential Road	0.095	0.077	0.093	0.261	0.219
Large Road	-0.057	-0.012	-0.038	-0.107	-0.103
Cycleway	-0.016	-0.019	-0.002†	-0.028	-0.016
Footway	0.987	0.947	0.891	2.419	1.991
Shared Path	0.109	0.062	0.079	0.226	0.196
Living Street	0.321	0.134	0.258	0.804	0.718
Pedestrian zone	0.478	0.351	0.416	1.280	1.080
Steps	23.725	21.180	18.235	60.126	48.112
Surface Type					
Unpaved	0.213	0.188	0.201	0.540	0.450
Cobble	0.252	0.153	0.111	0.478	0.396
Surf. Unknown	0.252	0.303	0.261	0.615	0.497
Wrong Way	0.678	0.672	0.656	2.038	1.629
N. parameters	**20**	**20**	**20**	**21**	**22**
Final LL	**-416,713.9**	**-416,628.4**	**-407,414.7**	**-405,157.4**	**-404,893.2**
Adj. ρ^2	**0.283**	**0.284**	**0.299**	**0.303**	**0.304**
BIC	**833,665**	**833,494**	**815,067**	**810,564**	**810,048**

Table 1 compares the route choice specifications. The baseline value-of-distance (VoD) model reaches an adjusted ρ^2 of 0.283. Replacing distance with cyclist-specific cruising time alone (VoT_{no}) produces almost identical fit, indicating that cruising speed by itself does not improve explanatory power relative to distance.

By contrast, introducing empirical travel-time costs from eq. (1) yields a marked improvement. The expected-time model (VoT_{exp}) raises the adjusted ρ^2 to 0.299 and increases log-likelihood by more than 9,000 units relative to VoD. Infrastructure and surface coefficients also decrease in magnitude, suggesting that part of their apparent effect in VoD models reflects underlying time losses due to local frictions. Route choice is therefore better explained when these time implications are introduced explicitly.

Further gains are obtained when expected travel time is added alongside distance-weighted infrastructure attributes: the VoD+t_{exp} model achieves an adjusted ρ^2 of 0.303. The best-performing linear specification decomposes travel time into cruising time and delay, yielding an adjusted ρ^2 of 0.304. Both components enter negatively and significantly, with delay weighted more strongly, consistent with findings for motorised travel, confirming the behavioural relevance of stop-and-go conditions.

Delays account for roughly 20% of total bicycle travel time in our dataset. If cyclists could cruise continuously, their trips would therefore be about one-fifth shorter. Combined with the higher marginal disutility of delay relative to cruising time, this highlights the importance of interruptions, waiting, and re-acceleration penalties in perceived route cost. It also reinforces the value of measures that reduce friction, such as wider cycling infrastructure or better signal coordination—in improving experienced route quality.

Taken together, the results show that explanatory gains arise not from modelling cruising speed in isolation but from capturing the full travel-time burden experienced by cyclists. The poor performance of VoT$_{no}$ compared with the much stronger performance of expected-time and delay-based models confirms that local slowdowns and waiting times are central to route choice behaviour.

Infrastructure coefficients remain broadly stable across model specifications. Steps, footways, pedestrian areas, wrong-way cycling, and poor surfaces carry substantial disutility, while separated tracks and painted lanes on larger roads are generally favoured. Several effects attenuate in the VoT$_{exp}$ model, supporting the interpretation that infrastructure preferences partly reflect unmodelled time penalties.

Logarithmic specifications were also estimated and provided slightly higher fit but did not alter the ranking of the route-cost formulations. For clarity, we therefore report only the linear results.

5 Discussion and future research

The results show that cyclists' route choices are better explained by calibrated travel-time measures than by distance-based infrastructure proxies alone. In particular, the relatively weak performance of the cruising-time-only specification, compared with the substantially stronger performance of the expected-time and delay-based models, indicates that delays arising from travel frictions are a central component of perceived route cost. This supports the view that many infrastructure attributes influence behaviour primarily through their effects on stop-and-go conditions and local travel-time losses. It also helps explain why the value-of-time bicycle route choice model in Liu et al. (2025) was outperformed by the corresponding value-of-distance specification: their predicted speeds and delays were neither cyclist-specific nor link-specific, and therefore could not capture the heterogeneous time implications revealed in our analysis.

A limitation of the present framework is that the delay term aggregates stopping, waiting, rolling slowdowns, and residual noise into a single component, which is sufficient for modelling expected travel time but less useful for interpreting underlying behavioural mechanisms or constructing richer generalised cost measures.

A natural direction for future work is to extend the framework to include energetic costs alongside travel time. Cycling effort depends not only on cruising speed and slope but also on repeated deceleration and re-acceleration at intersections and other friction points. Capturing these processes would require a more explicit treatment of intersection dynamics, including stopping probability, waiting time, and restart profiles. This, in turn, would benefit from richer data: higher-resolution speed trajectories around intersections, contextual information such as signal control and traffic conditions, rider- and bicycle-specific characteristics (e.g., mass, bicycle type), and environmental factors such as wind or more detailed elevation. More broadly, future research could incorporate cyclist heterogeneity more systematically and develop more transferable route-cost representations that remain rooted in observed riding behaviour.

References

Ben-Akiva, M., & Bierlaire, M. (1999). Discrete Choice Methods and their Applications to Short

Term Travel Decisions. In R. W. Hall (Ed.), *Handbook of Transportation Science* (pp. 5–33). Boston, MA: Springer US. Retrieved 2023-07-05, from https://doi.org/10.1007/978-1-4615-5203-1 doi: 10.1007/978-1-4615-5203-1_2

Bigazzi, A., & Lindsey, R. (2019, June). A utility-based bicycle speed choice model with time and energy factors. *Transportation*, *46*(3), 995–1009. Retrieved 2025-05-21, from https://doi.org/10.1007/s11116-018-9907-2 doi: 10.1007/s11116-018-9907-2

Broach, J., Dill, J., & Gliebe, J. (2012, December). 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. Retrieved 2025-11-21, from https://www.sciencedirect.com/science/article/pii/S0965856412001164 doi: 10.1016/j.tra.2012.07.005

Cazor, L., Duncan, L. C., Staves, C. E., Isenschmid, U., Paulsen, M., & Rasmussen, T. K. (n.d.). *Formulation of energy exertion in bicycle networks and estimation of its impact on route choices.*

Liu, C., Johansson, F., Kristoffersson, I., & Rydergren, C. (2025, June). Estimating bicycle route choice model in value-of-time space, a case study for Gothenburg. *Journal of Cycling and Micromobility Research*, *4*, 100064. Retrieved 2025-11-07, from https://www.sciencedirect.com/science/article/pii/S2950105925000087 doi: 10.1016/j.jcmr.2025.100064

Ma, X., & Luo, D. (2016, July). Modeling cyclist acceleration process for bicycle traffic simulation using naturalistic data. *Transportation Research Part F: Traffic Psychology and Behaviour*, *40*, 130–144. Retrieved 2025-05-22, from https://www.sciencedirect.com/science/article/pii/S1369847816300195 doi: 10.1016/j.trf.2016.04.009

Meister, A., Felder, M., Schmid, B., & Axhausen, K. W. (2023, July). Route choice modeling for cyclists on urban networks. *Transportation Research Part A: Policy and Practice*, *173*, 103723. Retrieved 2024-01-17, from https://www.sciencedirect.com/science/article/pii/S096585642300143X doi: 10.1016/j.tra.2023.103723

Newson, P., & Krumm, J. (2009, November). Hidden Markov map matching through noise and sparseness. In *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems* (pp. 336–343). Seattle Washington: ACM. Retrieved 2025-10-27, from https://dl.acm.org/doi/10.1145/1653771.1653818 doi: 10.1145/1653771.1653818

Łukawska, M. (2024, September). Quantitative modelling of cyclists' route choice behaviour on utilitarian trips based on GPS data: associated factors and behavioural implications. *Transport Reviews*, *44*(5), 1045–1076. Retrieved 2025-11-19, from https://doi.org/10.1080/01441647.2024.2355468 (_eprint: https://doi.org/10.1080/01441647.2024.2355468) doi: 10.1080/01441647.2024.2355468

Łukawska, M., Paulsen, M., Rasmussen, T. K., Jensen, A. F., & Nielsen, O. A. (2023, October). A joint bicycle route choice model for various cycling frequencies and trip distances based on a large crowdsourced GPS dataset. *Transportation Research Part A: Policy and Practice*, *176*, 103834. Retrieved 2023-11-15, from https://www.sciencedirect.com/science/article/pii/S0965856423002549 doi: 10.1016/j.tra.2023.103834