Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

# Bridging the freight data deficit: An integrated multi-resolution OD matrix inference architecture and survey shelf-life metric

**Wei Sun**

wei.sun@adelaide.edu.au

Adelaide University

**Akshay Vij**

Adelaide University

**Nicholas Kaliszewski**

Adelaide University

**Research Article**

**Keywords:**

**Posted Date:** September 29th, 2026

**DOI:** https://doi.org/10.21203/rs.3.rs-10512032/v1

**License:** This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

**Additional Declarations:** No competing interests reported.

# Bridging the freight data deficit: An integrated multi-resolution OD matrix inference architecture and survey shelf-life metric

Wei Sun (corresponding author)[1], Akshay Vij[2], Nicholas Kaliszewski[3]

[1,2,3]Adelaide University, Adelaide, Australia

Email: wei.sun@adelaide.edu.au

## Abstract

Transport authorities globally face a persistent road-freight data deficit, frequently forced to appraise multi-billion-dollar infrastructure using infrequent, aging legacy surveys and continuous, but origin-destination (OD)-blind, traffic counts. To bridge this gap, this paper proposes an integrated, multi-resolution OD-inference architecture designed for data-poor planning environments. Rather than calibrating counts directly against an outdated survey or relying on simple growth factors, the architecture embeds a conditionally projected structural prior within a scalable, single-level calibration workflow. This mitigates volume obsolescence and provides a contemporaneously informed macroeconomic anchor. To navigate the spatial scale mismatch between coarse reporting zones and granular network routing, the architecture utilizes an intermediate calibration geography. Because this intermediate layer inherently introduces weakly identified spatial degrees of freedom, an application-specific hierarchical regulariser is deployed to constrain unsupported local redistribution and enforce macroeconomic plausibility. Using the island state of Tasmania, Australia, as a controlled closed-network test bed, the architecture successfully generates plausibility-preserving 2024 OD estimates without the confounding influence of boundary flows. Finally, through a controlled longitudinal experiment (2015–2024) using a consistent sensor subset, we operationalise "survey shelf-life" by quantifying the structural divergence between the aging prior and current network reality. Ultimately, this research provides a transferable, extensible blueprint for continuous strategic OD inference and an evidence-based monitoring signal to guide capital-intensive survey investments.

## 1 Introduction

Effective appraisal of multi-billion-dollar road freight infrastructure requires comprehensive physical origin-destination (OD) matrices that capture the entirety of network usage. Under ideal conditions, inferring a high-fidelity, behaviourally plausible OD matrix relies on four components: (i) a contemporary, full-scope survey serving as a structural prior; (ii) empirical, high-penetration route-choice data (e.g., GPS trajectories); (iii) consistently maintained official freight network graphs; and (iv) contemporary traffic counts.

However, transport authorities globally operate within a severe data-deficit reality. Commercial freight is highly fragmented and logistics data are typically privatised, largely precluding the use of empirical route-choice data at a state or national scale. Official freight-specific network graphs are frequently unavailable, forcing reliance on open-source proxies. Most critically, as comprehensive freight surveys become increasingly cost-prohibitive and subject to declining response rates, planners are rarely equipped with contemporary structural data, relying instead on discontinued, aging legacy surveys (Holguín-Veras et al., 2011). Consequently, planners face an analytical compromise: they must either extrapolate a decaying structural prior or rely solely on sparse observational constraints (traffic counts) that cannot reliably answer OD-based strategic questions.

Methodologically, inferring OD matrices from link counts is an intrinsically under-determined inverse problem (Van Zuylen and Willumsen, 1980). Without a high-quality structural prior,

calibration algorithms are mathematically vulnerable to severe, weakly identified spatial redistribution, arbitrarily inflating local, short-distance trips simply to satisfy a traffic count. Conversely, purely structural projection suffers from cumulative temporal drift. While modern single-level optimisation architectures have successfully bypassed the computational burden of traditional bi-level equilibrium models (Spiess, 1990; Sun et al., 2023), their inference quality remains critically bottlenecked by the quality of the initial seed and their inherent vulnerability to spatial overfitting when empirical trajectory data are unavailable.

Addressing this planning problem requires more than applying a conventional prior-anchored OD estimator. First, a legacy survey should not be treated uncritically as a current calibration target: its aggregate volume and spatial-economic structure may have diverged from the contemporary freight system. A defensible workflow must therefore distinguish the role of the legacy survey as a source of structural knowledge from its role as a direct representation of current demand. Second, the spatial resolutions at which structural evidence, route choice, and traffic observations are available rarely coincide. Coarse reporting zones preserve interpretability but are too aggregated for reliable network assignment, whereas fine routing zones create an excessively under-determined calibration problem. These temporal and spatial mismatches motivate an integrated, multi-resolution inference architecture.

Acknowledging that the foundational mathematical mechanics of structural distribution and prior-anchoring are well-established (Cascetta, 1984; Van Zuylen and Willumsen, 1980), this study does not propose a new OD-estimation theory. Instead, it demonstrates an integrated, multi-resolution planning architecture for producing plausibility-preserving freight OD evidence under severe data constraints. The study makes three contributions:

1. **Integrated prior projection and count calibration.** We embed a conditionally projected structural prior within a count-based calibration workflow. Rather than calibrating directly against an aging survey matrix or relying only on simple growth adjustment, the workflow uses contemporary socio-economic covariates to mitigate volume obsolescence and refresh the macroeconomic basis of the prior. The resulting prior remains conditional on the stability of the estimated structural relationships; count calibration then reconciles it with contemporary observed network use.

2. **Multi-resolution spatial stabilisation.** We separate reporting, routing, and calibration geographies, using intermediate-resolution calibration and application-specific dual-scale regularisation to limit unsupported local redistribution.

3. **Longitudinal data-governance application.** We use divergence between the projected prior and calibrated OD estimate as a controlled survey shelf-life monitoring signal.

The remainder of this paper is organized as follows. Section 2 positions the work within the literature on structural freight OD projection and count-based inference. Section 3 details the proposed modular architecture. Section 4 describes the study area and data controls. Section 5 presents the Tasmania results for the calibrated 2024 ODM and the 2015–2024 longitudinal decay series. Section 6 discusses implications for planning practice, transferability, limitations, and finally, it concludes the study.

## 2 Literature review

Freight origin-destination matrices (ODMs) are foundational inputs for strategic infrastructure appraisal, freight planning, and network performance assessment. Yet, they remain among the most difficult transport quantities to observe directly at scale. Compared with passenger travel, freight OD inference is complicated by commercial data sensitivities, heterogeneous vehicle classes, supply-chain-driven spatial structures, empty repositioning, and complex

multi-stop delivery patterns. In practice, the central difficulty for transport agencies is rarely a complete absence of data, but rather a temporal and spatial mismatch between intermittent, information-rich structural sources such as freight surveys and continuous, but OD-blind, observational sources such as traffic counts. This challenge is especially acute in data-deficit planning environments, where comprehensive freight surveys are infrequent, discontinued, or prohibitively expensive.

### 2.1. Structural projection and its temporal limits

A long-standing approach to freight demand modelling is to synthesize a prior ODM from the spatial distribution of economic activity and land use. In transport practice, this is commonly operationalized through trip-end generation-attraction models combined with gravity-type trip distribution, where flows are determined by zonal mass variables such as freight-related employment and a deterrence term based on distance or travel time. Early gravity applications relied heavily on extensive survey data; however, later work demonstrated that gravity parameters could also be calibrated under assumed assignment conditions using traffic observations, thereby reducing dependence on costly OD surveys (Tamin and Willumsen, 1989).

Structural models remain attractive for strategic planning because they explicitly anchor demand to interpretable socio-economic drivers. Input-output and commodity-flow approaches provide an alternative structural paradigm by linking freight movements to inter-industry transactions and macroeconomic structure (Barna, 1963; Leontief, 1936; Pompigna and Mauro, 2020). These models are valuable for commodity-specific and long-term forecasting, but they do not always align cleanly with the practical notion of full-scope truck-trip ODMs required for highway planning. In particular, bottom-up commodity frameworks often emphasize major freight tasks and heavy commodities, while under-representing the broader physical road task associated with empty truck repositioning, urban service trips, and light commercial vehicle activity.

The critical limitation of purely structural approaches is temporal drift. As land use, industrial composition, logistics organization, and freight-routing practices evolve, legacy structural relationships progressively lose explanatory power. Consequently, ageing forward a historical prior without continuous empirical correction from the physical road network becomes increasingly unreliable. For planning agencies operating with a discontinued freight survey, the problem is therefore not simply how to extrapolate the prior, but how to preserve its macroeconomic interpretability while updating it to remain consistent with contemporary network use.

### 2.2. Count-based inference and under-identification

A second major convention approaches the problem from the network upward, seeking to infer an ODM that reproduces observed link or site flows when assigned to a transport network. The core difficulty is that OD inference from traffic counts is intrinsically under-determined: the number of unknown OD flows generally far exceeds the number of independent count constraints. The practical implication is that count observations do not identify a complete ODM on their own. Estimation therefore depends on the credibility of supplementary structural information and on explicit choices about how strongly the inferred matrix should remain anchored to that information. Prior-anchored and Bayesian formulations formalise this trade-off, but they do not remove the need to consider whether the prior remains temporally and spatially appropriate for the estimation year (Abrahamsson, 1998; Bell, 1991; Maher, 1983).

Classical approaches fall broadly into entropy-based and statistical families. Van Zuylen and Willumsen (1980) formulated a most-likely trip matrix approach using information minimisation and entropy principles. Cascetta (1984) and Bell (1991) developed generalized least squares

(GLS) formulations that balance count fit against deviation from a prior matrix, thereby allowing different confidence levels to be assigned to the prior and the observations. Subsequent work focused increasingly on computational tractability. Spiess (1990), for example, proposed a highly influential gradient-based OD adjustment method for iteratively correcting a starting matrix to better match observed traffic flows.

A further stream of research introduced route-choice endogeneity through bi-level formulations, coupling upper-level OD adjustment with lower-level network assignment, often under user-equilibrium assumptions (Yang et al., 1992). While theoretically appealing, such formulations become computationally burdensome for repeated longitudinal estimation over large regional systems. Moreover, freight routing does not always conform well to pure user-equilibrium logic, as commercial vehicles frequently operate under fixed schedules, delivery constraints, road-class restrictions, and corridor preferences that reflect logistics practice rather than generalized travel-time minimization alone (Hess et al., 2015). These limitations have motivated growing interest in more tractable single-level architectures that represent route choice through empirical or representative path sets rather than fully endogenous equilibrium assignment (Behara et al., 2022; Shen and Wynter, 2012; Sun et al., 2023).

However, replacing bi-level assignment with efficient single-level calibration does not eliminate the core inferential challenge. In sparse-count environments, calibration quality remains critically dependent on the quality of the initial structural seed and on how the framework manages under-identification across space. This challenge becomes more severe when legacy survey matrices are defined at coarse reporting geographies that are poorly aligned with the finer spatial detail needed for route generation and count matching.

Spatial resolution is itself a consequential modelling decision. Coarse zones preserve tractability and alignment with strategic reporting systems, but can obscure local access and route-assignment effects; highly disaggregated zones can improve spatial representation but substantially increase dimensionality and exacerbate under-identification when observation coverage is sparse. Transport-modelling research has consequently emphasised the importance of matching spatial resolution to modelling purpose and has developed multi-resolution approaches that allocate different levels of detail to different analytical functions (Hadi et al., 2022; Moeckel and Donnelly, 2015). These limitations make the quality, currency, and spatial interpretation of the calibration prior central to practical OD inference. The question is therefore not only how a prior should be constrained during calibration, but also how an aging prior should be prepared before calibration begins.

### 2.3. Updating legacy priors before calibration

A critical concern is that conventional count-based OD calibration commonly treats a historical survey or modelled matrix as a target prior, constraining the updated matrix to remain close to its historical structure while improving agreement with contemporary traffic counts. In some applications, the historical matrix is first adjusted using aggregate, zonal, or OD-specific growth factors. These approaches are practical and remain useful where the prior is recent and its spatial structure is considered broadly reliable. The broader OD-estimation literature has long recognised the need to update historical demand information using contemporary observations, including dynamic and real-time estimation of time-dependent OD flows (Ashok and Ben-Akiva, 2000; Cascetta et al., 1993). However, these approaches commonly address short-interval operational variation, whereas strategic freight planning must often contend with the longer-horizon obsolescence of a discontinued survey and the limited availability of contemporaneous OD observations.

However, direct use of an aging prior places two distinct burdens on the calibration stage. It must accommodate changes in overall demand magnitude while also redistributing flows in response to evolving spatial-economic and logistics conditions. In a sparse-count setting,

these adjustments are not independently identified: multiple combinations of volume and spatial redistribution can reproduce the same observed counts. Simple growth-factor adjustment may mitigate aggregate volume obsolescence, but generally preserves the historical spatial pattern and therefore cannot represent changes in the relationship between economic activity and freight movements.

An alternative is to treat a legacy ODM not as a current target matrix, but as evidence from which to estimate a structural prior-generation model. That model can then be conditionally projected using contemporary socio-economic covariates before count calibration. Such a projection does not recover the true contemporary ODM, as its validity depends on the temporal transferability of relationships estimated from the legacy survey, and it may mitigate rather than eliminate volume obsolescence. It can nevertheless provide a more contemporaneously informed and interpretable calibration anchor than direct use of an unadjusted historical matrix. The divergence between this conditional structural expectation and count-constrained estimates can then provide evidence about the continuing adequacy of legacy structural information.

### 2.4. Emerging trajectory data and multi-source fusion

Recent advances in Intelligent Transport Systems (ITS) have introduced additional data streams that partially observe vehicle trajectories and route choice, including GPS/telematics, Bluetooth or Wi-Fi re-identification, and Weigh-in-Motion (WIM) systems. These data can provide information on path use, corridor movements, and partial OD structure that conventional site counts alone do not directly observe. For example, sparse but strategically placed re-identification data can substantially reduce OD indeterminacy when fused with conventional count observations (Michau et al., 2019). Similarly, truck GPS trajectories have been used to infer destination choice behaviour and empirical transition structures (Basso et al., 2022; Demissie and Kattan, 2022).

Despite their value, these emerging data sources are rarely complete, stable, or representative at a network-wide planning scale. Freight telematics often have uneven penetration, are biased toward specific fleet operators, and may be commercially inaccessible to public agencies. Even where such data exist, they may not provide the longitudinal consistency or jurisdiction-wide coverage needed for routine strategic monitoring. This means that richer trajectory data should be viewed not as a universal replacement for traditional planning evidence, but as a supplementary information layer whose practical usefulness depends on coverage, representativeness, and institutional access. Their uneven availability reinforces the value of architectures that can operate with legacy surveys and routine counts as minimum inputs, while admitting richer trajectory observations where coverage and institutional access permit.

### 2.5. Integrated inference for data-deficit planning

These constraints have motivated integrated OD inference frameworks that combine a structural prior derived from surveys or land-use models with observational evidence from counts and, where available, partial trajectory data. In statistical terms, this can be interpreted as updating a prior belief using a likelihood derived from observed flows (Maher, 1983). In optimisation terms, it is a regularised calibration problem that minimizes count misfit while penalizing deviation from a prior matrix or from structure implied by the prior-generating model (Bell, 1991; Cascetta, 1984; Spiess, 1990).

Recent single-level data-fusion frameworks provide an operationally attractive basis for repeated regional estimation because they can incorporate representative route-choice structures and heterogeneous observations without solving a full equilibrium assignment within every OD update (Behara et al., 2022; Dey et al., 2020; Shen and Wynter, 2012; Sun

et al., 2023). In parallel, machine-learning and deep-learning approaches architectures (Abbas et al., 2024; Di et al., 2023; Peng et al., 2024; Yang et al., 2023) can learn complex OD-flow relationships where large, representative, and continuous trajectory datasets are available. Yet their use in data-deficit public-sector freight planning remains constrained by data availability, transferability, and the need for interpretable links between inferred demand and observed spatial-economic conditions.

Taken together, the literature suggests that neither structural projection nor observational calibration alone is sufficient in data-deficit freight-planning environments. Structural models preserve an interpretable macroeconomic basis for demand, but their conditional relationships may lose validity as logistics systems evolve. Count-based inference incorporates contemporary network observations, but remains under-identified without a credible prior. Richer trajectory data can reduce uncertainty, but may be incomplete, selectively available, or inaccessible to public agencies. These conditions motivate an integrated, multi-resolution inference architecture that: (i) learns structural information from a legacy survey and conditionally projects an interpretable prior using contemporary socio-economic data; (ii) reconciles that prior with contemporary heavy-vehicle observations through scalable count-based calibration; (iii) manages the mismatch between reporting, routing, and calibration geographies; and (iv) treats divergence between the projected prior and calibrated estimates as a monitoring signal for survey obsolescence and future data-investment decisions.

The contribution of the present study is not a new OD-estimation theory or a claim to recover unobserved OD truth. It is an empirical demonstration of how this integrated architecture can produce plausibility-preserving strategic freight evidence, stabilise a regional under-determined inference problem, and support evidence governance in a severe data-deficit setting.

# 3 Methodology

The proposed framework operationalizes a multi-resolution spatial architecture to navigate the fundamental trade-offs between Modifiable Areal Unit Problem (MAUP) distortions, computational tractability, and spatial overfitting in data-deficit settings.

The architecture relies on a nested geographic structure. While this paper instantiates the architecture using the Australian Statistical Geography Standard (ASGS) levels of SA1, SA2, and SA3, the underlying principle is fully transferable to any jurisdiction possessing hierarchically nested spatial data (e.g., census blocks nesting into traffic analysis zones or administrative districts). We define three distinct geographic tiers based on their analytical function:

- **SA3 (Reporting Geography):** The coarse resolution native to the legacy survey, preserving macroeconomic interpretability and stable socio-economic bounds.
- **SA1 (Routing Geography):** The highly granular resolution providing reliable network entry/exit points for generating behaviourally realistic route-choice sets.
- **SA2 (Calibration Geography):** An intermediate spatial resolution acting as the mathematical bridge for inference.

This three-tiered design is a motivated, pragmatic response to spatial-scale mismatch. While the legacy structural prior exists at the macro-level (SA3), directly routing and calibrating at this coarse level triggers severe MAUP distortions, as macro-centroids are too aggregated to reliably assign flows to physical traffic counts. Conversely, routing and calibrating entirely at the highly granular level (SA1) ensures network fidelity but results in computational

intractability and severe under-identification due to the extreme ratio of unknown OD pairs to sparse traffic counts.

Therefore, the architecture utilizes the intermediate level (SA2) to bridge the gap. We disaggregate the structural baseline to SA2 (Section 3.1) and aggregate the granular route choices to SA2 (Section 3.2). We deliberately avoid aggregating route choices all the way to SA3, as the intra-zonal proportionality assumptions required for such aggregation become highly unreliable at a macro scale. Finally, while SA2 strikes a necessary balance for computational tractability and routing fidelity, it still possesses vastly more degrees of freedom than the available count constraints. To mitigate this, we introduce an application-specific hierarchical regulariser at the SA2 level (Section 3.3) to act as a stabilisation mechanism, preventing the solver from engaging in unsupported spatial redistribution.

Let $Z_3$and $Z_2$denote the set of zones at the SA3 and SA2 levels, respectively. For any year $t$, the unknown target OD matrices are $Q_t^{(3)} = \{q_{ij,t}^{(3)}\}_{i,j \in Z_3}$ and $Q_t^{(2)} = \{q_{ab,t}^{(2)}\}_{a,b \in Z_2}$. The architecture fuses structural knowledge from a discontinued legacy survey (year $t_0$) with contemporary traffic counts $y_{s,t}$ at monitoring sites $s \in S$.

## 3.1 Structural learning and baseline projection

The structural projection module is explicitly designed as a model-agnostic architectural layer. Its primary function is to extract legacy spatial–economic patterns and project them forward using contemporary socio-demographics, thereby systematically mitigating volume decay and acting as a structural smoothing mechanism to provide a structurally informed initial seed and anchor for downstream calibration. Because the architecture is modular, this projection layer can seamlessly host highly sophisticated, non-linear projection engines—such as machine-learning architectures—in data-rich, large-scale environments (e.g., continental networks). However, for data-sparse regional applications and the present closed-system pilot, the legacy survey matrix may contain very few macro-reporting zones (e.g., 15 SA3 zones in Tasmania). To prevent the severe overfitting inherent to highly parameterized models operating on restricted sample sizes, we instantiate this module using a robust, classical two-step gravity framework.

**Trip Generation:** For each SA3 zone $i \in Z_3$, trip productions ($P_{i,t}$) and attractions ($A_{i,t}$) are modelled as regression functions of socio-demographic covariates $x_{i,t}$ (predominantly Freight-Related Employment, FRE):

$$P_{i,t_0} = g_P(x_{i,t_0}; \theta_P), A_{i,t_0} = g_A(x_{i,t_0}; \theta_A) \quad (1)$$

Zonal scaling factors (k-factors) capture unmodeled localized industrial effects and are carried forward deterministically.

**Trip Distribution:** Legacy OD flows are structurally mapped using a doubly constrained gravity model. Utilizing a power deterrence function $F(d;\beta) = d^{-\beta}$, the legacy flows satisfy:

$$q_{ij,t_0}^{(3)} = \alpha_i \gamma_j P_{i,t_0} A_{j,t_0} d_{ij,t_0}^{-\beta}, i,j \in Z_3 \quad (2)$$

Rather than utilizing dynamic, year-specific network travel times, we deliberately calculate impedance ($d_{ij}$) using a static SA1-adjacency graph distance proxy. This is a critical methodological control for longitudinal inference. By decoupling demand generation from volatile transport supply datasets (e.g., OpenStreetMap), we ensure that any observed longitudinal "structural decay" is driven mostly by demand-side economic drift, rather than spurious artifacts caused by year-to-year mapping inconsistencies.

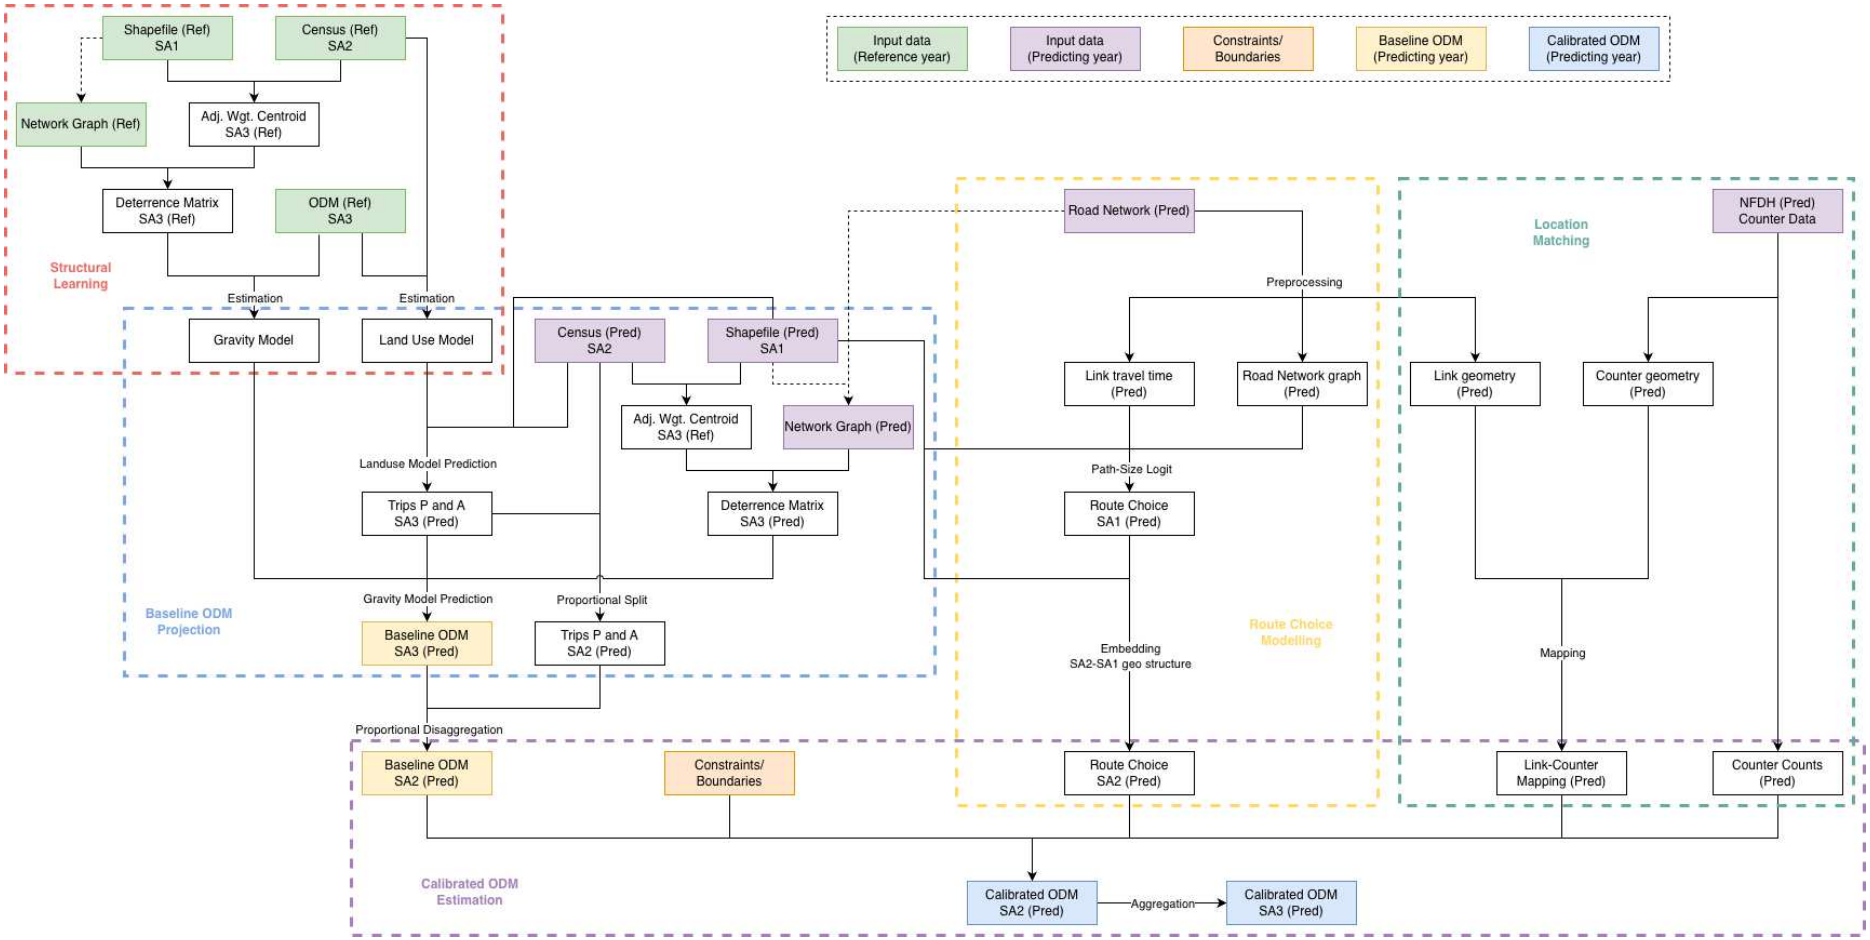
**Baseline Disaggregation:** Given target-year covariates $x_{i,t}$, we synthesize the projected macro-baseline $Q_t^{(3),\text{base}}$. Because direct calibration at SA3 is precluded by MAUP, we require a systematic rule to bridge the baseline to the SA2 calibration geography. We define a static, FRE-anchored proportional disaggregation rule. Let $m(a) \in Z_3$ map SA2 zone $a \in Z_2$ to its parent SA3 zone, and $w_a$ denote the static FRE weight. Within-SA3 production and attraction shares are:

$$\phi_a = \frac{w_a}{\sum_{a':m(a')=m(a)} w_{a'}}, \psi_b = \frac{w_b}{\sum_{b':m(b')=m(b)} w_{b'}} \quad (3)$$

For any SA3 matrix $Q^{(3)}$, its proportional SA2 disaggregation is strictly defined as:

$$q_{ab}^{(2),\text{dis}}\left(Q^{(3)}\right) = q_{m(a)m(b)}^{(3)} \phi_a \psi_b \quad (4)$$

**Figure 1- Architecture overview (Australian data context)**



## 3.2 Single-level route choice and observation mapping

Direct path generation between macro-zones is behaviourally degenerate. We deploy a two-stage aggregation protocol to connect unobserved OD demand to physical counts.

First, a Path-Size Logit (PSL) model generates behaviourally diverse route sets $K_{rs}$ at the fine SA1 level $(r, s \in Z_1)$. Unlike passenger vehicles, commercial heavy-vehicle routing is governed by logistics cost-minimization, fixed delivery windows, and rigorous road-class accessibility restrictions, rendering it highly deterministic and corridor-bound (Behara et al., 2022; Hess et al., 2015). Consequently, utilizing a stochastic PSL model over a freight-filtered OpenStreetMap network provides a systematic and defensible approximation of physical routing in the absence of empirical GPS trajectories. The utility for path $k$ is:

$$V_{k,rs} = \beta_c C_{k,rs} + \beta_t T_{k,rs} + \beta_{ps} \ln(\mathrm{PS}_{k,rs}) \quad (5)$$

In the absence of local empirical trajectory data to estimate jurisdiction-specific coefficients, the path utility was parameterized to reflect real-world commercial freight behaviour using official economic guidelines. The time and cost coefficients $(\beta_t, \beta_c)$ were derived from the Transport for NSW Economic Parameter Values (TfNSW, 2025), representing contemporary heavy-vehicle operating costs and value of time, indexed to June 2024 prices. The path-size coefficient $(\beta_p)$ was adopted from established heavy-vehicle routing literature (e.g., Hess et al., 2015) to penalise overlapping routes and prevent the artificial inflation of choice probabilities for heavily correlated paths.

Follow-up calculation yields the logit path probability:

$$p_{k|rs} = \frac{\exp(V_{k,rs})}{\sum_{k' \in K_{rs}} \exp(V_{k',rs})} \quad (6)$$

Defining $\delta_{\ell k} \in \{0,1\}$ as the link incidence indicator, the SA1 OD-specific link utilization probability is:

$$\pi_{\ell|rs} = \sum_{k \in K_{rs}} p_{k|rs}\, \delta_{\ell k} \quad (7)$$

Second, to achieve computational tractability, SA1 probabilities are aggregated to the SA2 calibration level. Letting $\mathcal{R}_a$ denote the SA1 sub-zones within SA2 zone $a$, the SA2 link-use probability is derived via intra-SA2 uniform weighting $\omega_{rs}^{(ab)}$:

$$\pi_{\ell|ab} = \sum_{r \in \mathcal{R}_a} \sum_{s \in \mathcal{R}_b} \omega_{rs}^{(ab)}\, \pi_{\ell|rs}, \text{ where } \sum_{r \in \mathcal{R}_a} \sum_{s \in \mathcal{R}_b} \omega_{rs}^{(ab)} = 1 \quad (8)$$

The predicted flow at monitoring site $s$, mapping to network links $\ell \in \mathcal{M}(s)$, is the linear combination of the unknown SA2 matrix and the static routing probabilities:

$$f_{s,t}\left(Q_t^{(2)}\right) = \sum_{\ell \in \mathcal{M}(s)} x_{\ell,t}\left(Q_t^{(2)}\right), \text{ where } x_{\ell,t}\left(Q_t^{(2)}\right) = \sum_{a,b \in Z_2} q_{ab,t}^{(2)}\, \pi_{\ell|ab} \quad (9)$$

### 3.3 Calibrated ODM estimation and hierarchical regularisation

The purpose of the hierarchical regulariser is not to claim a new general class of OD regularisation. Rather, it is designed to resolve an identifiability consequence of the proposed multi-resolution architecture. Calibration at SA2 is necessary to improve the spatial correspondence between route assignment and count locations, but this intermediate geography introduces substantially more OD degrees of freedom than can be identified from sparse traffic counts alone. The dual-scale penalty therefore acts as a structural stabilisation mechanism: it permits count-informed adjustment of macro-level flows while discouraging unsupported within-macro redistribution at the intermediate level.

Define the aggregation operator $A_{3\leftarrow 2}$ mapping SA2 flows to SA3 totals:

$$q_{ij,t}^{(3)}(Q_t^{(2)}) = \sum_{a:m(a)=i} \sum_{b:m(b)=j} q_{ab,t}^{(2)}, i,j \in Z_3(10)$$

Rather than anchoring SA2 flows to a static historical prior, we construct an endogenous "dynamic alternative" SA2 ODM. This acts as the micro-spatial stabiliser, representing what the flows would look like if they adhered strictly to localized FRE proportions (Eq. 4), updating dynamically based on the solver's current SA3 macro-inference:

$$Q_t^{(3)}\left(Q_t^{(2)}\right) = A_{3\leftarrow 2}\left(Q_t^{(2)}\right) \quad (11)$$

$$Q_t^{(2),\text{alt}}\left(Q_t^{(2)}\right) = q^{(2),\text{dis}}\left(Q_t^{(3)}\left(Q_t^{(2)}\right)\right) \quad (12)$$

Let $D(\cdot,\cdot)$ denote structural cosine distance. The single-level nonlinear programming (NLP) calibration is formulated as:

$$\min_{Q_t^{(2)}\geq 0} \lambda_3 D\left(A_{3\leftarrow 2}(Q_t^{(2)}), Q_t^{(3),\text{base}}\right) + \lambda_2 D\left(Q_t^{(2)}, Q_t^{(2),\text{alt}}(Q_t^{(2)})\right) \quad (13)$$

Subject to hard count and systemic volume constraints:

$$(1-\epsilon_s)y_{s,t} \leq f_{s,t}(Q_t^{(2)}) \leq (1+\epsilon_s)y_{s,t}, \forall s \in S$$

$$(1-\epsilon_T)T_t^{\text{base}} \leq \sum_{a,b\in Z_2} q_{ab,t}^{(2)} \leq (1+\epsilon_T)T_t^{\text{base}}$$

where $T_t^{\text{base}}$ is the projected baseline total volume, and $\epsilon_s, \epsilon_T$ define tight feasibility relaxations.

This dual-penalty objective function provides the structural stabilization: the macro-regulariser ($\lambda_3$) enforces adherence to the regional structural baseline, while the endogenous micro-regulariser ($\lambda_2$) explicitly penalizes deviations from proportional localized land-use. Unlike classical entropy maximization (Van Zuylen and Willumsen, 1980), which penalizes deviations from a static, single-resolution historical prior, this multi-resolution architecture allows macro-flows to flexibly calibrate to observed network usage while preserving proportional land-use logic at the micro-level.

In this data-deficit prototype, the regularisation weights are set equally (a 1:1 ratio). In the absence of full-scope empirical ground truth to definitively prioritize the reliability of the macro-economics over the micro-proportions, equal weighting serves as an unbiased heuristic. The 1:1 ratio is a prototype design choice, not an empirically calibrated parameter; the present sensitivity analysis (see footnote) tests local robustness, not global optimality. As richer trajectory data becomes available in future deployments, these weights can be empirically optimized using L-curve analysis[1].

Furthermore, we explicitly acknowledge that this formulation constitutes a non-linear and non-convex programming (NLP) problem, primarily due to the endogenous nature of the dynamic alternative (Eq. 12) coupled with the cosine distance metric. Because theoretical convergence to a global optimum cannot be strictly guaranteed, the problem was solved using the Ipopt interior-point algorithm via the Pyomo Python library, supplemented by a multi-start heuristic protocol. By introducing random scalar perturbations (±10%) to the baseline seed, we observed that the solver consistently converged to highly structurally similar matrices, ensuring the primary planning output (Eq. 14) remains behaviourally stable and robust against initial conditions:

$$Q_t^{(3),\text{cal}} = A_{3\leftarrow 2}\left(Q_t^{(2)}\right) \quad (14)$$

### 3.4 Structural decay monitoring

The architecture uniquely isolates longitudinal structural decay (spatial redistribution) from simple volume decay (handled deterministically by baseline FRE scaling and bounded by $\epsilon_T$). Using scale-invariant cosine similarity, the "survey shelf-life" deviation metric is formally defined as:

$$\text{decay}(t) = 1 - \text{cosim}\left(Q_t^{(3),\text{cal}}, Q_t^{(3),\text{base}}\right) \quad (15)$$

By tracking this metric over time against a constant sensor subset, we systematically quantify the progressive degradation of the legacy prior's explanatory power.

## 4 Data and experimental controls

To empirically validate the architecture, the island state of Tasmania, Australia, was utilized as a closed-system test bed. The architecture itself is designed to be fully transferable to any region possessing a legacy structural prior, contemporary socio-economic covariates, a routable network, and stable traffic counts. However, scaling to open continental networks introduces the massive confounding variables of porous land borders, unobserved boundary flows, and "through-traffic," which typically require specialized gateway-hub modelling. By contrast, Tasmania's island geography naturally eliminates these externalities. This provides a pristine, controlled environment to isolate the fundamental stabilisation logic, multi-resolution calibration behaviour, and longitudinal drift mechanics of the inference framework prior to continental deployment.

The foundational structural prior was derived from the discontinued 2014 Australian Bureau of Statistics (ABS) Road Freight Movements Survey (RFMS). Representing the last comprehensive, full-scope snapshot of Tasmanian road freight, the 2014 SA3-level sub-matrix

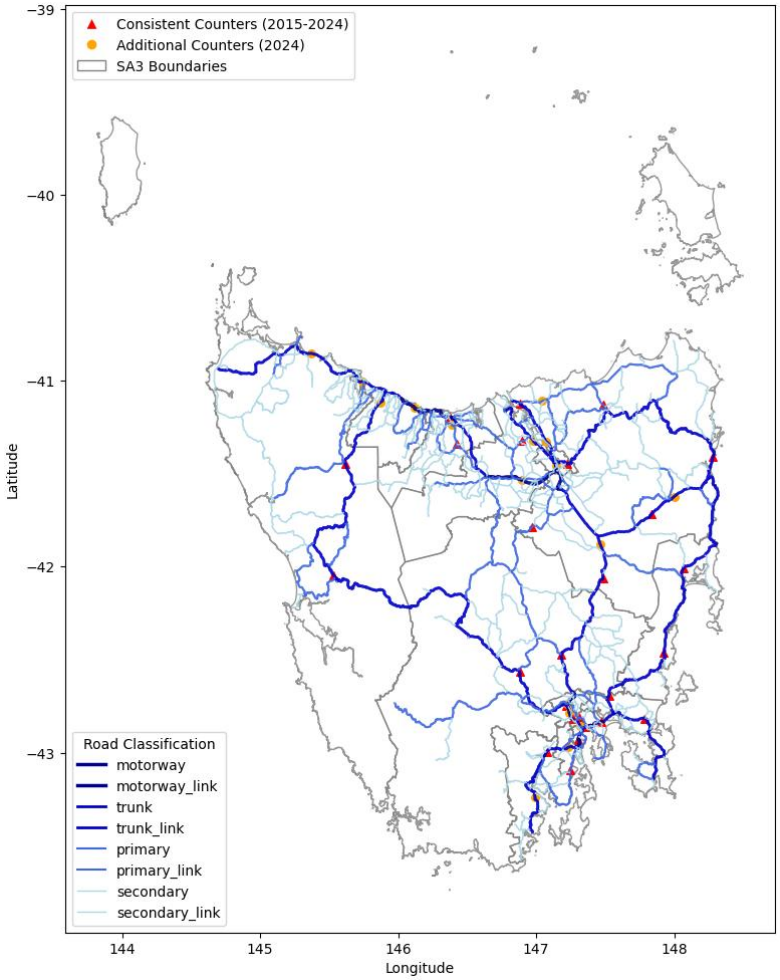
[1] We have tested 1:2 and 2:1 ratios, and the resulting calibrated ODM remained highly structurally similar with cosine similarity above 0.98. However, we acknowledge this establishes local structural robustness under nearby ratios, not parameter optimality.

was not treated as a static "current" calibration target, but rather as the vital training dataset for the macro-level spatial-economic structural projection module.

To drive the baseline projection across the 2015–2024 monitoring period, socio-demographic covariates were extracted from ABS Census datasets (2011, 2016, 2021) and extrapolated to the target years. The primary generation driver was Freight-Related Employment (FRE), calculated by aggregating employment data across core supply-chain industries (e.g., agriculture, manufacturing, mining, retail, and transport/warehousing).

To prevent topological inconsistencies between demographic data and the physical road network, a strict spatial pre-processing step was applied. Disconnected offshore SA1 zones (e.g., King Island, Flinders Island) were systematically excluded from the SA1 adjacency graph, as their freight movements are strictly maritime/aviation and do not interact with the main island road network. Consequently, an “SA2 connected ratio” was computed to proportionally adjust zonal FRE totals prior to aggregation. This ensures that the demand generation drivers accurately reflect only the economic activity serviceable by the road network.

**Figure 2-Tasmanian SA3 boundaries, strategic freight road networks and traffic counter locations**




The empirical calibration constraints ($y_{s,t}$) consisted of heavy-vehicle traffic counts provided by permanent monitoring sites managed by the Tasmanian Department of State Growth. Because historical metadata detailing lane-specific directionality was insufficient, observations were conservatively treated at the bi-directional site level. To ensure rigorous longitudinal

analysis, the deployment of these constraints was strictly bifurcated based on the inference objective:

1. **Point-in-time calibration (2024):** Utilized all 42 available active counters to maximize observational fidelity and spatial constraint for the target year.

2. **Longitudinal decay monitoring (2015–2024):** Strictly restricted to a consistent subset of 26 permanent counters operational across the entire ten-year period. This vital control isolates demand-side structural drift; dynamically adding new counters mid-series would artificially tighten the optimization constraints, fatally skewing the temporal deviation metric and invalidating the longitudinal shelf-life experiment.

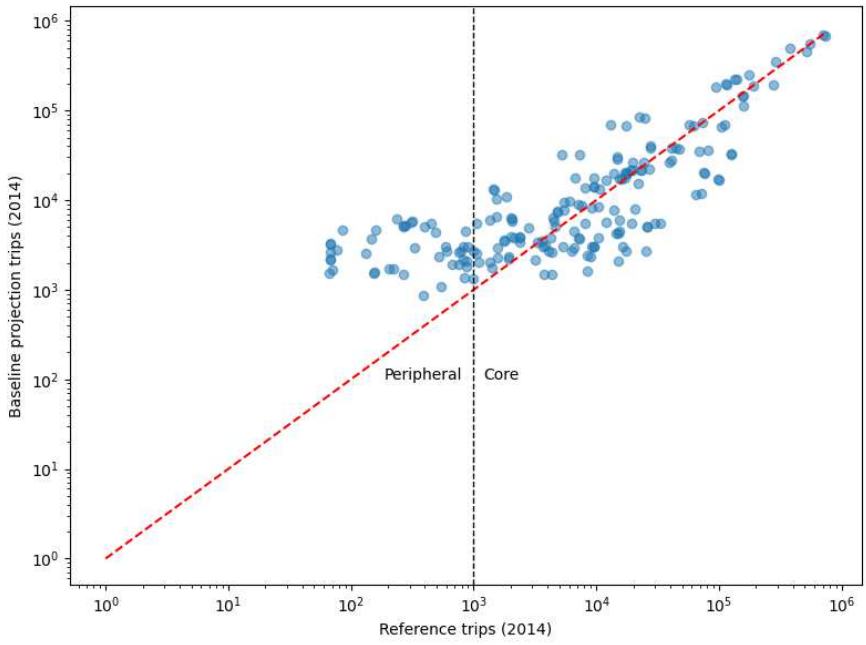
# 5 Results and analysis

## 5.1 Validation of the structural prior and routing substrate

Prior to target-year inference, the structural learning and routing modules were strictly validated against the 2014 RFMS to ensure the framework captured the underlying spatial–economic logic. The Path-Size Logit (PSL) model successfully mapped behaviourally diverse route options, utilizing 91% of the strategic network graph without pathological concentration. The gravity deterrence parameter converged optimally at $\beta = 1.6310$ ( $R^2 = 0.928$ ), representing a realistic behavioural composite of regional long-haul and urban service vehicles.

To assess if the projected baseline is robust for future projection, the synthesized 2014 macro-baseline was benchmarked against the raw survey ODM **(Figure 3)**. The overall cosine similarity between the two matrices was highly aligned at 0.96. Crucially, 94% of total modelled freight trips occurred in OD pairs where the model's prediction was statistically indistinguishable from the survey (falling within the survey's native Relative Standard Error confidence intervals). This confirms that the structural model acts as a robust "information filter"—stripping out random, non-systematic sampling noise and discreteness, while retaining the foundational macroeconomic corridors required for stable future-year projections.

**Figure 3 - Scatter plot of Survey vs Baseline ODMs (2014)**



## 5.2 Point-in-time inference and the ablation study (2024)

The 2024 calibration fused the projected economic baseline with 42 contemporary traffic constraints. To facilitate convergence while preserving the macroeconomic integrity of the prior, the total system volume constraint ($\epsilon_T$) was configured to a tight 0.01 relaxation. The solver shifted total systemic travel marginally by ~1%, confirming that the optimization engine satisfied local counts purely via spatial OD redistribution rather than artificially inflating overall macroeconomic demand. Internal sanity checks demonstrated near-perfect bidirectional flow symmetry ($r = 0.999$), confirming the solver did not generate highly directional, pathological flows simply to match localized counts. Furthermore, Z-score analysis confirmed that calibrated zonal trip productions remained within historical uncertainty bounds, barring two zones reflecting decadal industrial shifts.

**Figure 4-SA3 trip distribution heatmaps for 2024 (Structural Baseline vs Traffic-Calibrated)**




This spatial redistribution is visualized in Figure 4, which compares the 2024 unconstrained baseline projection against the final calibrated ODM at the macro-reporting (SA3) level. The fundamental spatial pattern derived from the 2014 economic logic, specifically the heavy intra-zonal flows along the main diagonal and the primary inter-regional corridors, remains clearly visible. However, the calibration successfully introduces vital localized adjustments, intensifying certain corridors and dampening others to reconcile the theoretical land-use projection with the physical reality of 2024 network usage. This spatial shift (yielding a final SA3 cosine similarity of ≈0.88) highlights the exact value of the hybrid inference step: it captures physical routing evolutions that land-use covariates alone cannot predict.

Achieving this visually plausible result in an under-determined system is entirely dependent on the proposed hierarchical architecture. To empirically demonstrate this mathematical necessity, an ablation study was conducted by removing the lower-level micro-spatial anchor (setting $\lambda_2 = 0$). The results were stark. Without the SA2 regulariser, the macro-level SA3 cosine similarity artificially inflated to 0.958, while the underlying SA2 structure collapsed to a similarity of just 0.132. Rather than proving that hierarchical regularisation is universally necessary for all OD inference, this divergence illustrates its necessity specifically for this intermediate-SA2 architecture given sparse counts and no trajectory data. Without a micro-spatial stabiliser, the solver engages in severe, weakly identified spatial redistribution. To preserve the macro-level SA3 totals while satisfying local traffic counts, the unconstrained

solver generates highly unstable local SA2 routing that is entirely inconsistent with the specified land-use-based micro-spatial reference structure. By contrast, the full hierarchical formulation forces a balanced, behaviourally plausible compromise, maintaining an SA2 similarity of ≈0.937 alongside the SA3 stability. Ultimately, the hierarchical regulariser acts as a vital structural stabiliser, ensuring the inferred matrix respects both macro-economic volume and micro-spatial land-use reality.

Cross-sectional validation against independent, bottom-up commodity datasets yielded weak or non-existent rank correlations (Tasmanian Freight Survey $\rho = 0.344$; CSIRO TraNSIT $\rho = 0.068$). This low correlation highlights a well-known divergence in freight modelling: bottom-up commodity models excel at capturing heavy, selective commodity flows (~8–11% of the total task), whereas physical traffic counters, and therefore our calibrated ODM, capture the holistic network reality, including empty truck repositioning, urban logistics, and light commercial vehicles. However, we acknowledge that this discrepancy may also reflect unquantified assumptions in our static routing proxy, warranting future investigation using empirical GPS data. Consequently, in a severe data-deficit environment, deterministic cell-by-cell validation is precluded; the calibrated 2024 ODM represents a highly constrained, regularised best-estimate that successfully absorbs modern traffic realities without violating fundamental economic geography.

## 5.3 Longitudinal structural decay: Operationalizing "Survey shelf-life"

Leveraging the strict experimental controls (constant network topology and a consistent 26-counter subset), the framework successfully isolated economic drift from artifacts of changing sensor density. Figures 5 and 6 track the structural distance (cosine distance) between the theoretically projected baseline and the observationally calibrated reality at both the macro planning level (SA3) and the micro routing level (SA2) across a decade.

**Figure 5-SA3 level longitudinal structural decay (2015-2024)**

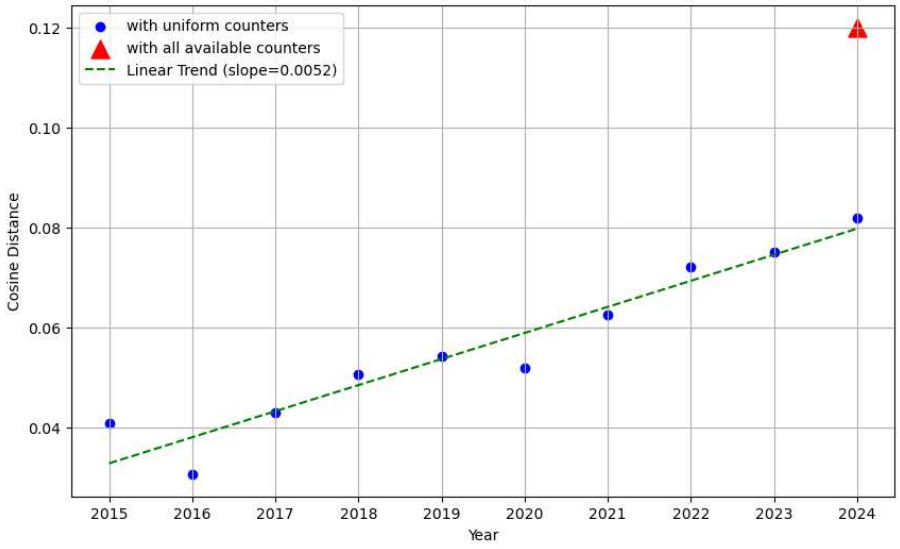
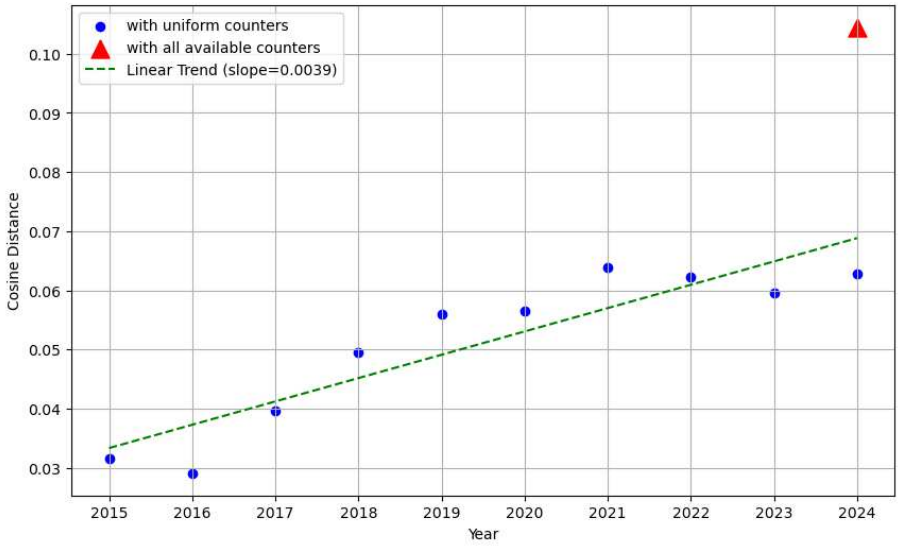



**Figure 6-SA2 level longitudinal structural decay (2015-2024)**




At the macro-reporting level (Figure 5), the time-series exhibits a clear, almost monotonic increasing trend, rising from a cosine distance of ~0.04 in 2015 to over 0.08 by 2024. This confirms that the explanatory power of the 2014-derived structural knowledge is progressively degrading. Notably, a slight plateau occurs during the COVID-19 pandemic (2020), which aligns with the unique resilience of road freight: as initial transport restrictions froze discretionary economic activity, the network temporarily reverted to its most foundational macroeconomic corridors. At the micro-spatial routing level (Figure 6), the overarching decadal trend mirrors the macro-decay but exhibits noticeably higher year-to-year volatility. This variance is a mathematical inevitability of the data-deficit environment: possessing vastly more degrees of freedom relative to the sparse counts, the SA2 matrix inherently absorbs the bulk of the system's temporary observational noise, localized routing detours, and transient supply-chain shocks.

Crucially, Figure 5 reveals the strategic value of observational density. The red triangle represents the structural distance for 2024 when all 42 available counters are utilized. This value (0.12) is significantly higher than the 26-counter trendline (0.08). The constraint of sparse networks allows under-determined solutions to stay comfortably, yet artificially, close to the prior. Increasing counter density forces the model to confront more empirical reality, dramatically accelerating the early detection of structural anomalies. By tracking this controlled trajectory, transport authorities can define an evidence-informed monitoring threshold, calibrated to local decision needs, to trigger the commissioning of a new, capital-intensive freight survey, removing the reliance on arbitrary fixed schedules.

# 6 Discussion and conclusions

This study addresses the pervasive data deficit that characterizes modern strategic freight planning. By integrating a conditionally projected legacy survey with continuous, OD-blind traffic counts via a multi-resolution, single-level calibration architecture, we successfully generated an internally consistent, updated 2024 ODM for the closed-system test bed of Tasmania, Australia. This architecture effectively bridges the gap between purely structural models, which suffer from cumulative temporal drift, and conventional count-based calibration, which is mathematically under-identified and spatially unstable in sparse environments.

A key design feature of this architecture is its multi-resolution spatial structure supported by hierarchical regularisation. In severely under-determined environments, this dual-scale penalty acts as a vital structural stabiliser. By anchoring micro-level spatial proportions to localized land-use variables while granting macro-level flows the flexibility to calibrate to

observed physical realities, the solver is prevented from engaging in weakly identified, extreme spatial redistribution. Consequently, the architecture provides a plausibility-preserving estimate of the aggregate heavy-vehicle network task, avoiding the potentially significant spatial overfitting that plagues classical OD estimation when high-fidelity trajectory data are unavailable.

Strategically, this research operationalizes the concept of "survey shelf-life" into a quantifiable policy metric. By conducting a controlled longitudinal experiment, we isolated the systematic structural decay of the 2014 legacy prior over a decade. For transport authorities, this transitions the question of survey-refresh timing from an arbitrary political cycle to an evidence-based monitoring problem. Importantly, a deviation threshold crossing (e.g., a cosine distance of 0.10) cannot independently trigger a major survey investment. Rather, the metric provides a monitoring signal that agencies must combine with locally specified decision thresholds, uncertainty tolerances, and the expected value of new survey information. Furthermore, the analysis reframes the value of continuous network monitoring: sparse count networks allow models to remain artificially close to outdated priors, whereas denser permanent counter deployments rigorously accelerate the detection of true structural drift, acting as critical strategic assets for evidence governance.

Despite these advancements, fundamental limitations inherent to data-deficit inference must be acknowledged. First, the unobservability of deterministic ground truth means these inferred matrices represent a highly constrained, regularised state rather than absolute cell-by-cell precision; they capture aggregate heavy-vehicle flows but cannot independently identify empty repositioning or discrete multi-stop tours. Second, the reliance on a static routing network isolates demand-side drift but deliberately omits supply-demand co-evolution over the decade. Furthermore, the generalized Path-Size Logit model abstracts away operational realities such as peak-hour congestion, vehicle-class accessibility restrictions, and dynamic toll avoidance. Finally, the heuristic 1:1 weighting of the hierarchical regulariser serves as a neutral assumption of ignorance in the absence of full-scope ground truth; future applications with richer data should formalize this via empirical sensitivity analysis (e.g., L-curve cross-validation).

The core contribution of this study is the architectural workflow itself, which is intentionally modular and extensible. Its transferability to other jurisdictions depends on four preconditions: (i) a legacy structural prior, (ii) temporally updated spatial-economic covariates, (iii) a routable road network, and (iv) a stable subset of traffic observations. While Tasmania provided an ideal, controlled closed-network test bed for isolating internal inference stability, scaling this architecture to open continental networks (e.g., mainland Australia or Europe) requires the explicit treatment of unobserved boundary flows and through-traffic via gateway or gate-hub formulations. Furthermore, extreme spatial sparsity at a continental scale may necessitate replacing the classical gravity baseline with non-linear machine learning generators (e.g., hurdle XGBoost architectures). Ultimately, this research establishes a replicable, scientifically rigorous transitional pipeline. It provides transport agencies globally with a practical tool to maintain plausibility-preserving strategic freight evidence under severe data constraints, while establishing an architectural scaffold ready to ingest richer Intelligent Transport Systems (ITS) data as regional data infrastructure matures.

## Acknowledgements

This study was funded by the Australian Government Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts (DITRDCSA) and the iMOVE Cooperative Research Centre (CRC). The project was managed by David Mitchell and Joe O’Sullivan, Bureau of Infrastructure and Transport Research Economics (BITRE).

## References


Abbas, S., Sampedro, G.A., Almadhor, A., Abisado, M., Marzougui, M., Kim, T., Alasiry, A., 2024. PODE: privacy-enhanced distributed federated learning approach for origin-destination estimation. PeerJ Comput. Sci. 10, e2050. https://doi.org/10.7717/peerj-cs.2050

Abrahamsson, T., 1998. Estimation of Origin-Destination Matrices Using Traffic Counts- A Literature Survey. Work. Pap., Working Papers.

Ashok, K., Ben-Akiva, M.E., 2000. Alternative Approaches for Real-Time Estimation and Prediction of Time-Dependent Origin–Destination Flows. Transp. Sci. 34, 21–36. https://doi.org/10.1287/trsc.34.1.21.12282

Barna, T. (Ed.), 1963. Structural Interdependence and Economic Development. Palgrave Macmillan UK, London. https://doi.org/10.1007/978-1-349-81634-7

Basso, F., Pezoa, R., Tapia, N., Varas, M., 2022. Estimation of the Origin-Destination Matrix for Trucks That Use Highways: A Case Study in Chile. Sustainability 14, 2645. https://doi.org/10.3390/su14052645

Behara, K.N.S., Bhaskar, A., Chung, E., 2022. Single-level approach to estimate origin-destination matrix: exploiting turning proportions and partial OD flows. Transp. Lett. 14, 721–732. https://doi.org/10.1080/19427867.2021.1932182

Bell, M.G.H., 1991. The estimation of origin-destination matrices by constrained generalised least squares. Transp. Res. Part B Methodol. 25, 13–22. https://doi.org/10.1016/0191-2615(91)90010-G

Cascetta, E., 1984. Estimation of trip matrices from traffic counts and survey data: A generalized least squares estimator. Transp. Res. Part B Methodol. 18, 289–299. https://doi.org/10.1016/0191-2615(84)90012-2

Cascetta, E., Inaudi, D., Marquis, G., 1993. Dynamic Estimators of Origin-Destination Matrices Using Traffic Counts. Transp. Sci. 27, 363–373. https://doi.org/10.1287/trsc.27.4.363

Demissie, M.G., Kattan, L., 2022. Estimation of truck origin-destination flows using GPS data. Transp. Res. Part E Logist. Transp. Rev. 159, 102621. https://doi.org/10.1016/j.tre.2022.102621

Dey, S., Winter, S., Tomko, M., 2020. Origin–Destination Flow Estimation from Link Count Data Only. Sensors 20, 5226. https://doi.org/10.3390/s20185226

Di, X., Shi, R., Mo, Z., Fu, Y., 2023. Physics-Informed Deep Learning for Traffic State Estimation: A Survey and the Outlook. Algorithms 16, 305. https://doi.org/10.3390/a16060305

Hadi, M., Zhou, X., Hale, D., 2022. Multiresolution Modeling for Traffic Analysis: Guidebook [WWW Document]. URL https://rosap.ntl.bts.gov (accessed 7.23.26).

Hess, S., Quddus, M., Rieser-Schüssler, N., Daly, A., 2015. Developing advanced route choice models for heavy goods vehicles using GPS data. Transp. Res. Part E Logist. Transp. Rev. 77, 29–44. https://doi.org/10.1016/j.tre.2015.01.010

Holguín-Veras, J., Jaller, M., Destro, L., Ban, X. (Jeff), Lawson, C., Levinson, H.S., 2011. Freight Generation, Freight Trip Generation, and Perils of Using Constant Trip Rates. Transp. Res. Rec. J. Transp. Res. Board 2224, 68–81. https://doi.org/10.3141/2224-09

Leontief, W.W., 1936. Quantitative Input and Output Relations in the Economic Systems of the United States. Rev. Econ. Stat. 18, 105. https://doi.org/10.2307/1927837

Maher, M.J., 1983. Inferences on trip matrices from observations on link volumes: A Bayesian statistical approach. Transp. Res. Part B Methodol. 17, 435–447. https://doi.org/10.1016/0191-2615(83)90030-9

Michau, G., Pustelnik, N., Borgnat, P., Abry, P., Bhaskar, A., Chung, E., 2019. Combining traffic counts and Bluetooth data for link-origin-destination matrix estimation in large urban networks: The Brisbane case study. https://doi.org/10.48550/ARXIV.1907.07495

Moeckel, R., Donnelly, R., 2015. Gradual rasterization: redefining spatial resolution in transport modelling. Environ. Plan. B Plan. Des. 42, 888–903. https://doi.org/10.1068/b130199p

Peng, T., Gan, M., Ou, Q., Yang, X., Wei, L., Rødal Ler, H., Yu, H., 2024. Railway cold chain freight demand forecasting with graph neural networks: A novel GraphARMA-GRU model. Expert Syst. Appl. 255, 124693. https://doi.org/10.1016/j.eswa.2024.124693

Pompigna, A., Mauro, R., 2020. Input/Output models for freight transport demand: a macro approach to traffic analysis for a freight corridor. Arch. Transp. 54, 21–42. https://doi.org/10.5604/01.3001.0014.2729

Shen, W., Wynter, L., 2012. A new one-level convex optimization approach for estimating origin–destination demand. Transp. Res. Part B Methodol. 46, 1535–1555. https://doi.org/10.1016/j.trb.2012.07.005

Spiess, H., 1990. A gradient approach for the OD matrix adjustment problem.

Sun, W., Vij, A., Kaliszewski, N., 2023. A flexible and scalable single-level framework for OD matrix inference using IoT data. Transp. Res. Part Policy Pract. 175, 103775. https://doi.org/10.1016/j.tra.2023.103775

Tamin, O.Z., Willumsen, L.G., 1989. Transport demand model estimation from traffic counts. Transportation 16. https://doi.org/10.1007/BF00223044

TfNSW, 2025. Transport for NSW Economic Parameter Values. Transport for NSW, Sydney, Australia.

Van Zuylen, H.J., Willumsen, L.G., 1980. The most likely trip matrix estimated from traffic counts. Transp. Res. Part B Methodol. 14, 281–293. https://doi.org/10.1016/0191-2615(80)90008-9

Yang, H., Sasaki, T., Iida, Y., Asakura, Y., 1992. Estimation of origin-destination matrices from link traffic counts on congested networks. Transp. Res. Part B Methodol. 26, 417–434. https://doi.org/10.1016/0191-2615(92)90008-K

Yang, Y., Jia, B., Yan, X.-Y., Chen, Y., Song, D., Zhi, D., Wang, Y., Gao, Z., 2023. Estimating intercity heavy truck mobility flows using the deep gravity framework. Transp. Res. Part E Logist. Transp. Rev. 179, 103320. https://doi.org/10.1016/j.tre.2023.103320