Quantifying uncertainty in travel demand forecasts[1]

Dr Stuart Donovan (presenter), Senior Fellow, Motu Research stuart.donovan@motu.org.nz
Peter Clark, Director, UrbanUmber

ABSTRACT

Transport planning and investment processes often use models to generate travel demand forecasts that extend decades into the future. Currently, most transport models treat the future as being deterministic when it is, in practice, inherently uncertain. In this paper, we outline an approach to quantifying the effects of uncertainty in transport model inputs, such as changes in population and employment, on travel demand forecasts. By combining Bayesian regression models with Monte Carlo simulation methods, the approach enables us to quantify the effects of uncertainty in inputs while remaining computationally tractable. To demonstrate the approach, we present illustrative results for a proposed transport project in New Zealand, namely Auckland Light Rail. Although this application is only illustrative and further technical development is needed, we suggest the approach could – even now and in its current state – help to quantify uncertainty in travel demand forecasts and support better transport planning and investment decisions.

[1] The authors wish to acknowledge this research is published with the kind permission of Auckland Light Rail Limited. We extend our thanks to Tommy Parker, whose support made this work possible; The New Zealand Infrastructure Commission Te Waihanga for reviewing and refining earlier versions of this paper; Luis "Pilo" Willumsen for his encouragement and invaluable insights; and Simon Nielsen and Collins Teye of Transport for London and Elliot Clayton from Invise Limited for their insights. All remaining errors are our own.

INTRODUCTION

Transport planning and investment processes often rely on travel demand forecasts that extend decades into the future. These forecasts are commonly derived from travel demand models that make deterministic assumptions about the future, such as changes in population and employment. In practice, and as many researchers have noted, the future is inherently uncertain. Willumsen (2014), for example, observes “uncertainty is unavoidable in any forecast of future conditions”.

In response to uncertainty in future conditions, current transport planning and investment processes regularly make use of sensitivity testing. These tests proceed by allowing key inputs and assumptions in travel demand models to vary. Although such tests can reveal information on the relative sensitivity of outcomes to inputs and assumptions, they are often undertaken in a relatively ad-hoc manner. Sensitivity testing requires practitioners, for example, to identify the relevant inputs that are to be varied – often in isolation – as well as the magnitude of the variation to test.

In this paper, we present an alternative approach to quantifying uncertainty in travel demand forecasts that has the potential to be more systematic. This approach combines Bayesian regression models with Monte Carlo simulation methods to provide richer insights into uncertainty, while remaining computationally tractable. Although further development is needed – for example, empirical analyses of historical data to identify appropriate probability distributions and parameters to use in the Monte Carlo simulations – we suggest the approach is, even in its current state, sufficiently well-developed that it could inform transport planning and investment processes.

The following sections of this paper are structured as follows: First, we summarise the planning context and provide a background to travel demand modelling; second, we discuss the methodology for our proposed approach, which combines Bayesian regression models with Monte Carlo simulation methods; and third, we apply the approach to a case study of a project from Auckland, New Zealand, namely Auckland Light Rail. To finish, we discuss implications of our research; note several limitations of the approach; and consider opportunities for further research.

BACKGROUND

The planning context

In New Zealand and many other countries, planning and investment in major transport projects and policies usually proceed via structured business case processes.

Due to the long lifetimes and large capital costs associated with many transport projects and policies, these business case processes usually rely heavily on transport models that generate travel demand forecasts extending decades into the future. Forecasts from travel demand models are, for example, often used as inputs into economic appraisal processes, such as cost-benefit analysis (CBA), which seek to quantify economic effects over periods of 40 years, if not longer.[2]

To arrive at these forecasts, travel demand models in turn rely on inputs and assumptions – such as, for example, on population and employment growth – which are inherently uncertain. As Willumsen & Ortúzar (2016) note, “the future is not deterministic and no amount of technical skills and models can change that”. Uncertainty is especially relevant for major projects which have long planning horizons and where infrastructure lasts for many decades. Changes in demographics, land use, technology, and travel behaviour can, over such long timelines, all have large effects on travel demands. In this context, there is growing interest in methods that quantify the effects of uncertainty on the outcomes of transport planning processes.

Travel demand modelling

The travel demand forecasts generated by transport models play a key role in transport policy and

[2] In related research, for example, Lieswyn (2011) considers the effects of uncertainty on cost-benefit analyses of a potential transport investment in Nelson, New Zealand.

investment decision-making. These models seek to predict how people will travel in the future, based on, for example, socio-economic, demographic, and network conditions. Over decades of development, transport models have evolved into relatively sophisticated tools (McNally, 2007).

In many jurisdictions, so-called four-step models (FSMs) remain the dominant approach to travel demand modelling. The “four-steps” in FSMs can be summarised as:

1. *Trip generation*, which predicts how many trips originate in specific zones based on, for example, population, employment, and other aspects of the prevailing land uses.
2. *Trip distribution*, which determines the destination of trips generated in each zone over potential zones in response to, e.g. overall transport costs (“gravity model”).
3. *Mode choice*, which allocates trips travelling from one zone to another to the available transport modes (e.g., car, public transit, walking, cycling).
4. *Trip assignment*, which assigns trips from one zone to another by a specific mode to the underlying network, such as road links and public transport stops and services.

While FSMs have proven useful for modelling future transport outcomes, they are relatively resource intensive to operate. Running large-scale models requires not only significant computational resources but also careful management and oversight. Developing a specification, refining inputs, running simulations, and validating outputs can take weeks – particularly when testing multiple scenarios and/or generating outputs to meet specific business case needs.

More recently, a new class of transport models known as activity-based models (“ABMs”) have emerged and are being used in some jurisdictions. In New Zealand, for example, the Ministry of Transport has developed an ABM called “Monty”. In contrast to the four steps outlined above (that is, generation, distribution, mode choice and assignment), ABMs seek to predict the travel demands of individuals and households over the course of a typical day. Although their underlying mathematical structure differs from FSMs, ABMs rely on similar inputs and assumptions and are also computationally intensive. Hence, while we frame the approach presented in this paper in the context of conventional FSMs, such as that used to inform business vases in Auckland, we consider that the underlying motivations and methods are also relevant and transferable to ABMs.

Typologies of uncertainty

Uncertainty in travel demand forecasting arises from multiple sources, which we suggest can be broadly categorised into the following four typologies:

1. *Behavioural uncertainty*, which arises from uncertainty in the structure of behavioural models (e.g., the steps in FSMs vis-à-vis the individual approach of ABMs).
2. *Parameter uncertainty*, which arises from uncertainty in the estimated values of the parameters in the behavioural models (e.g., value of time and scoring penalties etc).
3. *Input uncertainty*, which arises from uncertainty in the current or future state of the world (e.g., socio-economics, demographics, networks, policies, and technologies).

Behavioural and parameter uncertainty – that is, the first two typologies noted above, or what Willumsen & Ortúzar (2016) describes as “model quality” – would seem to be most usefully explored during the development, calibration, and validation of travel demand models. That is, behavioural and parameter uncertainty should ideally be addressed *before* a travel demand model is used to inform planning and investment processes.

In contrast, input uncertainty arises primarily from the application of travel demand models to specific projects or policies. As the most relevant sources of uncertainty will vary widely between individual projects, input uncertainty is difficult to address ex ante and must instead be mostly tackled within individual business cases.[3] For this reason, we see input uncertainty as somewhat

[3] Input uncertainty can arise due to socioeconomic shocks, such as COVID and new technologies, which change the state of the world in which the broader transport system operates. Even though these shocks may be common between

distinct from the other two typologies of uncertainty noted above. Our focus on input uncertainty does not imply that behavioural and parameter uncertainty are unimportant but rather notes the latter's effects are best addressed during model development, calibration, and validation.

Moreover, input uncertainty afflicts both FSMs and ABMs via its effects on:

- *Socioeconomic and demographic outcomes*, such as population and employment;
- *Transport network attributes*, such as capacity and availability of infrastructure;
- *Transport policy settings*, such as tolls, fares, and parking policies; and
- *Technological changes*, such as work-from-home technologies.

Willumsen & Ortúzar (2016) posit that uncertainty associated with future data and scenarios – or, what we describe as input uncertainty – is likely to increase more rapidly in the future than uncertainty associated with model quality – or, what we have called structural and parameter uncertainty. For this reason, Willumsen & Ortúzar (2016) conclude it is especially important to adopt probabilistic approaches to address input uncertainty within business cases.

METHODOLOGY

In theory, one might potentially tackle input uncertainty directly by repeatedly sampling inputs and re-running transport models, whether they be FSMs or ABM. In practice, the number of potential input assumptions to test as well as the time involved in running travel demand models mean that such direct approaches are currently computationally infeasible.

The need for a computationally tractable approach to characterising input uncertainty is the primary motivation for this research. In this section, we present our approach for quantifying the effects of input uncertainty on travel demand forecasts. First, we describe the quantitative methods that we use and, second, we describe the case study that we use to illustrate our proposed approach.

Approach

Our approach hews relatively closely to that outlined by Willumsen (2014). Specifically, we combine Bayesian regression models with Monte Carlo simulation methods in two main steps:

- *Step 1 – Bayesian regression model.* In this step, we first identify a set of input variables that are considered relevant to uncertainty (for example, based on historical evidence, ex-post assessments, expert judgement, and stakeholder engagement). Second, we vary these input assumptions and re-run Auckland's travel demand model to quantify how changes in these inputs affect travel demands. Third, we use a Bayesian regression model to estimate the effects of changes in input assumptions on travel demand forecasts.
- *Step 2 – Monte Carlo simulation.* In this step, we first assume probability distributions and associated parameter values for the inputs that are of most interest to our case study. Second, we undertake a Monte Carlo simulation by repeatedly sampling sub-sets of model inputs from these distributions. Each individual sub-set of model inputs is then combined with random draws from the posterior distributions of the parameters estimated in step 1. This allows us to generate a distribution of travel demand forecasts in each future year.

Below, we introduce the case study that we use to illustrate the application of these two steps in our proposed approach to quantifying input uncertainty.

projects, the extent of their influence will still be quite project specific. For example, we can expect that uncertainty over the future uptake of work-from-home will be more relevant to major transport projects in larger urban centres.

Illustrative case study

To demonstrate our proposed approach, we use Auckland Light Rail (ALR) as an illustrative case study. By way of background, ALR is a proposed rapid transit project connecting Auckland's city centre to the Airport located to the south (Auckland Light Rail, 2024). Business cases for ALR relied on the Macro Strategic Model (MSM), which is a FSM operated by the Auckland Forecasting Centre (AFC). The MSM is regularly used to forecast travel demands in Auckland. During the two-year development of the ALR business case, MSM was used to model a total of 147 scenarios which provides us with a relatively rich source of data to illustrate the proposed approach.

We apply our approach to two illustrative scenarios for ALR:

- First, the "Baseline scenario" focuses on quantifying input uncertainty. To do so, we align assumptions with the original ALR business case and demonstrate how the proposed approach can be used to quantify uncertainty around the deterministic ("central") forecasts that were generated by the MSM and subsequently used in business cases.
- Second, the "Alternative scenario" allows input assumptions to differ from those used in the original ALR business case. Specifically, we assume demand starts low and rises over time. The aim of this scenario is to show how our proposed approach can be used in a "sketch model" mode to quickly quantify the effects of changes in inputs.

We use forecasts from MSM to explore and illustrate the merits of our proposed approach for quantifying input uncertainty, rather than to evaluate the ALR project itself. We emphasise that neither of the above two scenarios represent a formal review or critique of the original ALR business case, nor the assumptions, inputs, or views of any sponsoring agency.

RESULTS

Step 1 – Bayesian regression model

Identifying relevant variables

While multiple measures of passenger demand could be used as a dependent variable, the results presented here use passengers carried per hour per direction (p/h/d) at the peak load point (PLP) in the AM peak.[4] For the ALR project, stakeholders deemed this measure to be an important input into several transport planning and investment decisions, such as vehicle capacities, fleet requirements, and service frequency. Alternative demand measures can readily be used, however.

During the ALR business case process, the following input variables into the MSM were identified by stakeholders as both being most likely to influence passenger demand *and* being associated with relatively high levels of future uncertainty:

- *Population in the corridor*, which is influenced by a combination of market forces and central and local government housing policies, including but not limited to significant tracts of government-owned land along the corridor.
- *Employment in the city centre*, which is largely determined by local, national, and international economic market forces and planning regulations.
- *Work from Home (WfH) uptake*, which depends on both technology and industry and service sector employment policies.
- *Congestion pricing schemes*, which depends on future decisions made by both central and local government.

[4] The selection and definition of the variables used in the model, such as passengers at PLP in the AM peak and definition of the relevant corridor, introduces an additional potential source of uncertainty. We thank a reviewer for noting that the methodology we use to quantify uncertainty is also subject to uncertainty. In ou review, these sources of uncertainty will be explicitly accounted for in the uncertainty associated with our parameter estimates.

- *City centre parking prices*, which are influenced by both market forces and national and local government policies regarding the supply of parking.

Several other key drivers of demand, such as train frequency, capacity, and travel times, were identified and included in the regression model as they were important to its predictive power. As these drivers were deemed to lie more within the control of the operator or authority managing ALR – at least in the medium to long-run – they were not allowed to vary the Monte Carlo simulations undertaken in Step 2. Put simply, the uncertainty associated with these variables was assessed to be of somewhat reduced importance than the five explanatory variables listed above.

Estimating the regression model

We estimate a log-log specification as it performs well and allows for several of the resulting parameter estimates to be interpreted as constant elasticities, which in turn allows them to be compared to the wider transport economic literature. In our preferred log-log model, continuous variables were log-transformed, whereas dummy variables and those expressed as percentages were left untransformed. Our preferred specification for the linear model is as follows:

$$\log D_i = \beta_0 + \beta_1 \log P_i + \beta_2 \log E_i + \beta_3 \log I_i + \beta_4 \log C_i + \beta_5 \log F_i + \beta_6 W_i + \beta_7 \delta_i^C + \beta_8 \delta_i^P$$

Where for each of the 147 MSM runs denoted by i:

- $\log D_i$ denotes the natural log of passengers per hour per direction (p/h/d) at the peak load point (PLP) in the AM peak
- $\log P_i$ denotes the natural log of population in the corridor (000s)
- $\log E_i$ denotes the natural log of employment in the city centre (000s)
- $\log I_i$ denotes the natural log of in-vehicle travel-time by transit from the city centre to the airport, i.e. for the full length of the corridor (mins)
- $\log C_i$ denotes the natural log of total capacity at the peak load point (p/h/d, 000s)
- $\log F_i$ denotes the natural log of frequency (trains per hour)[5]
- W_i denotes work from home uptake (%), which is set to either 7% or 14%
- δ_i^C denotes a dummy for congestion charging (0=No, 1=Yes)
- δ_i^P denotes a dummy for parking prices (0=No, 1=Yes)
- β denote parameters in the regression model to be estimated.

Data on all variables is sourced from the MSM for future years of 2031, 2041, 2051, and 2065. Regression results for our preferred specification are shown in Table 1. For the log-transformed variables that can be interpreted as constant elasticities we find:

- *Corridor Population:* The estimate of 0.55 indicates a 10% increase in corridor population leads approximately to a 5.5% increase in peak load demand.
- *City Centre Employment:* The estimate of 0.89 indicates a 10% increase in City Centre employment results in an approximately 8.9% increase in peak load demand.
- In-Vehicle Travel Time: The estimate of -0.85 indicates a 10% increase in the in-vehicle travel time of ALR reduces peak load demand by approximately 8.5%.
- *Frequency.* The estimate of 0.16 indicates a 10% increase in frequency increased peak load demand by approximately 1.6%.

[5] Total capacity, C_i, and frequency, F_i, will be strongly positively correlated, because $C_i = V_i \cdot F_i$, where V_i = the capacity of rolling stock. To remove this correlation, future applications should model V_i and F_i instead of C_i and F_i.

Table 1: Regression results

Parameter	Variable	Interpretation / Measure	Estimate	s.e.	t-statistic
β_1	$\log P_i$	Population in corridor (000s)	0.55	0.14	4.01
β_2	$\log E_i$	Employment in city centre (000s)	0.89	0.21	4.33
β_3	$\log I_i$	In-vehicle travel time (mins)	-0.85	0.06	-15.26
β_4	$\log C_i$	Capacity at PLP (p/h/d, 000s)	0.14	0.01	9.43
β_5	$\log F_i$	Frequency (trains per hour)	0.16	0.04	3.81
β_6	W_i	WfH (7%, 14%)	-2.85	0.45	-6.27
β_7	δ_i^C	Congestion charge (0=No, 1=Yes)	0.02	0.02	0.62
β_8	δ_i^P	Parking price increase (0=No; 1=Yes)	0.13	0.02	5.33
R^2			*0.912*		
n (MSM model runs)			*147*		

The elasticity for in-vehicle time (-0.855) appears plausible when compared to external evidence. It is, for example, slightly higher than the upper end of observed short-term elasticities but close to long-term elasticities (Wallis, 2004; Balcombe, et al., 2004). Although the estimated elasticity for frequency (0.16) appears slightly lower than most estimates, this may also reflect that we are using peak load at the peak point as our demand measure – noting that most evidence finds larger frequency elasticities in off-peak periods. The relatively large t-statistics for estimates of continuous parameters indicate they are statistically significant at conventional levels and precisely estimated.

For non-transformed variables, the parameters can be interpreted as follows:

- *Work from Home:* The estimate of -2.85 suggests that an 1% increase in the proportion of people working from home reduces peak load demand by approximately 2.85%.
- *Parking Prices:* The estimate of 0.13 suggests that doubling parking prices in the city centre (from 0 to 1) increases peak load demand by approximately 13%.

Unfortunately, the estimated parameter for congestion pricing is small and imprecise, likely due to limited variation across the 147 scenarios. We return to this issue below when discussing the potential application of fractional factorial design to increase variation between scenarios.

In Figure 1, the left panel shows observed versus predicted values, where the vertical error bands indicate the 95% credibility intervals for the latter. For this case study, our preferred regression model has an R^2 of 0.912. This suggests the regression model explains most of the variance in peak load demand predicted by the MSM.

The right panel of Figure 1 then shows correlations between posterior estimates of the regression parameters that are of primary interest. We observe correlations with relatively large magnitudes between several parameter estimates, such as corridor population and city centre employment (-0.582) and capacity and frequency (-0.546). These correlations imply that when one parameter estimate is higher than average, the other will be lower than average, and vice versa. When we randomly sampling from these parameter estimates in the Monte Carlo simulation in Step 2, we can thus readily account for correlations between our parameter estimates.

Figure 1: Observed versus predicted values (left panel) and correlations between parameter estimates (right panel)

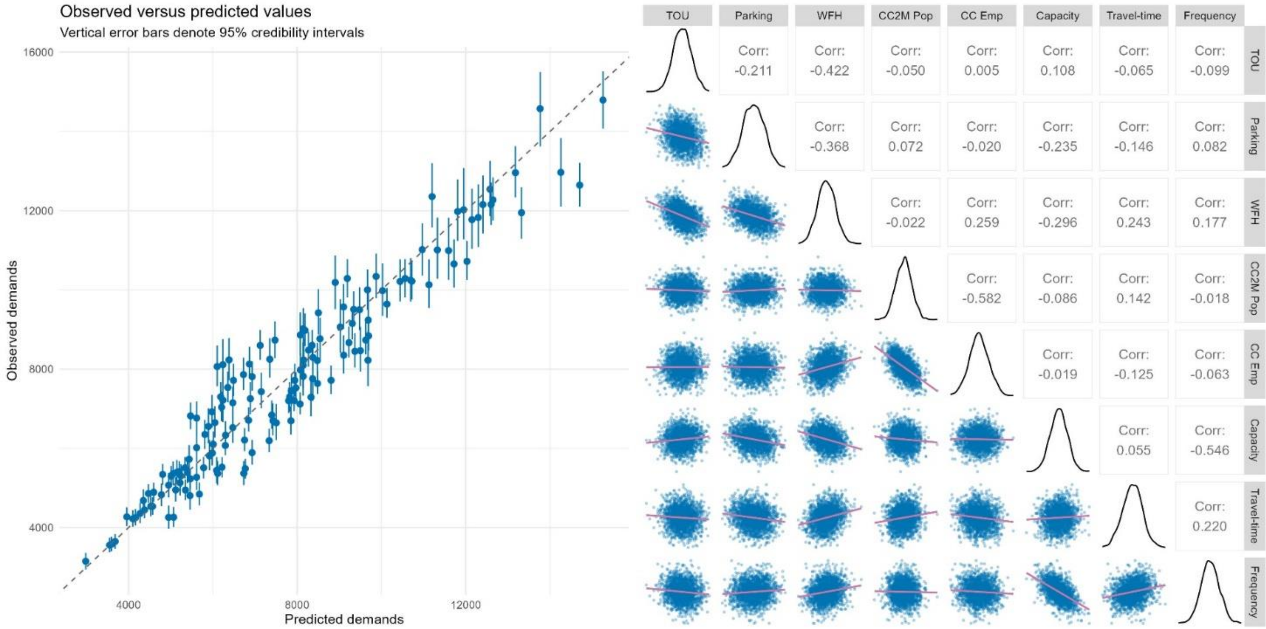

Notes: TOU denotes "time-of-use pricing"

Step 2 – Monte Carlo simulation

We now present the results of Monte Carlo simulations for the Baseline and Alternative scenarios described above: First, the "Baseline Scenario" adopts similar assumptions to travel demand modelling used in the business case for ALR and, second, the "Alternative Scenario", sets key inputs at low levels initially and then assumes they rise over time. Assumptions in both scenarios – such as statistical distributions and parameter values – are available on request from the authors.

In Figure 2, the shaded areas in the left and right panels illustrate the results of the Monte Carlo simulation, where the associated MSM forecasts in each year are denoted by crosses. The shaded intervals denote uncertainty as measured by the 80%, 90%, and 95% credibility intervals of estimates. These bands highlight the distribution of possible demand outcomes across simulations, providing insight into the variability and risk associated with projections.[6]

For the Baseline scenario, our proposed method produces outcomes that are very similar to those from the MSM. The 95% intervals approach approximately 10% of the forecast demand by 2066. In the Alternative scenario, demand initially starts below the MSM forecasts but grows more rapidly over time. This trend is consistent with the design of this scenario. Interestingly, even with what are quite different assumptions we find MSM forecasts lie within the 95% CI over this period.

[6] A reviewer noted there appears to be more downside risk than upside risk. This potentially hints at the effects of non-linearities in the underlying processes within the MSM that generate the data.

Figure 2: Results for the Baseline (left panel) and Alternative (right panel) scenarios

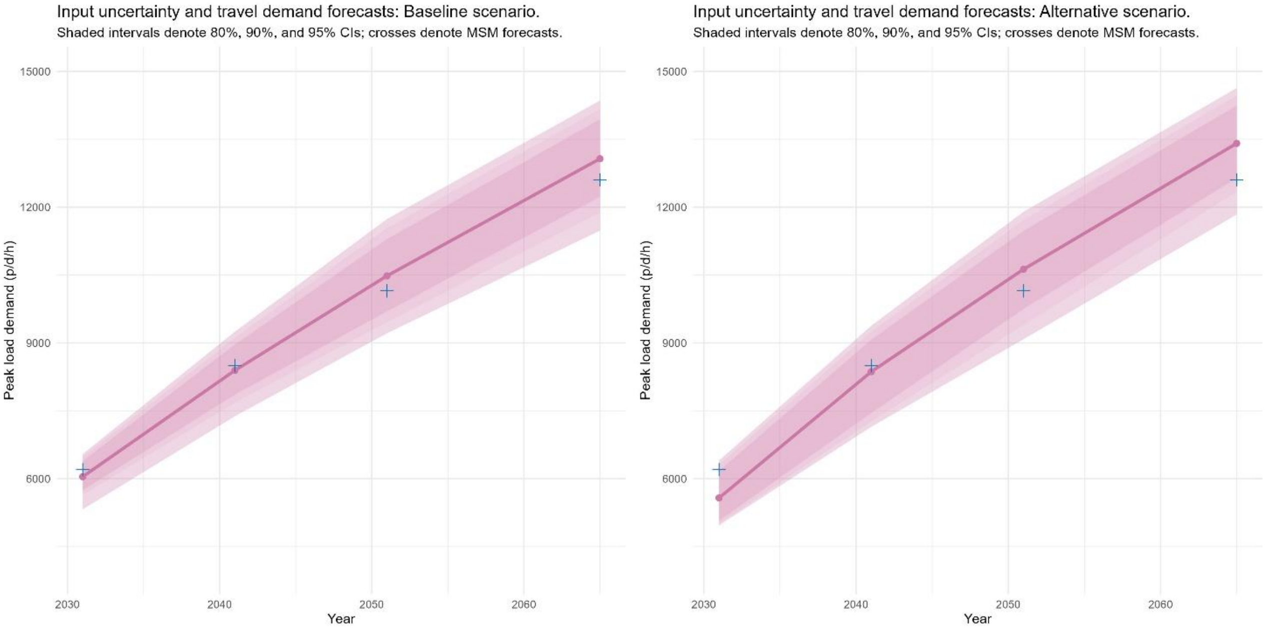

DISCUSSION, LIMITATIONS, AND FURTHER RESEARCH

Discussion

Our approach combines Bayesian regression models with Monte Carlo simulation methods to quickly quantify the uncertainty introduced by changes in key inputs into travel demand models.

For the Baseline scenario study above, uncertainty in selected inputs causes estimated demand in 2065 to vary by +/- 10% from the MSM forecasts. Although this sounds modest, we emphasise that it is the uncertainty associated with a small selection of inputs and does not include the effects of behavioural and parameter uncertainty arising from the travel demand model itself. As such, the actual range of uncertainty in travel demand forecasts is likely to be much larger than that shown above. Nonetheless, even allowing for uncertainty in a small subset of key inputs reveals considerable levels of uncertainty. We suggest this is an important point in itself.

Our proposed approach imposes relatively small additional computational demands. Both the Bayesian regression model and Monte Carlo simulation, for example, run in less than 5-minutes, even when using 10,000 samples. In contrast, current experience suggests it would be infeasible to run either FSMs or ABMs for this number of iterations. Of course, we had access to 147 runs of the MSM for this analysis, which took time and were generated over the course of a lengthy business case process.[7] Although this process may not be "typical", there are opportunities for more scientific approaches to reduce the number of runs that are required.

Limitations

We note three key limitations with our proposed approach:

- *Extensive model run requirements* – The Bayesian regression model was estimated on the results of 147 MSM transport model runs that were conducted over two years. This level of effort and timeframe may be unrealistic for many smaller projects.[8]

[7] We note that these runs were originally undertaken to inform the business case process rather than for the purposes of this analysis. In this context, the model run time can, in this instance, essentially be seen as a sunk cost of the business case process, rather than a marginal cost that is imposed by our analyses.
[8] Notwithstanding this limitation, we note that this level of effort might not be required for all projects. Rather, it could be that the application of this methodology to major projects might help to generate rules-of-thumb that could be applied more generally to other projects.

- *Limited variation in variables* – The MSM model runs undertaken to support the ALR business case resulted in limited variation in some variables, such as congestion pricing, which prevented the effects of these variables from being identified.
- *Detailed empirical analyses to support assumptions* – The Monte Carlo simulation requires assumptions for statistical distributions and associated parameters. Ideally, these assumptions would be carefully grounded in detailed empirical analyses.

While our case study serves to highlight the potential of this approach, its practical usefulness hinges on addressing these limitations. We consider these questions in the next section.

Further research

We note three main opportunities for further research. First, MSM runs for the ALR were not initially undertaken in a scientific way that sought to maximise variance and minimise multicollinearity. In Figure 4, we plot the correlation between attributes in our regression model. Here, blue indicates positive correlations, while red indicates negative relationships. High correlation values (close to ±1) suggest strong relationships, while values near 0 indicate weak or no association. Statistically significant correlations (1% level) are marked with a cross.

Figure 3: Correlation Plot of Regression Model Attributes

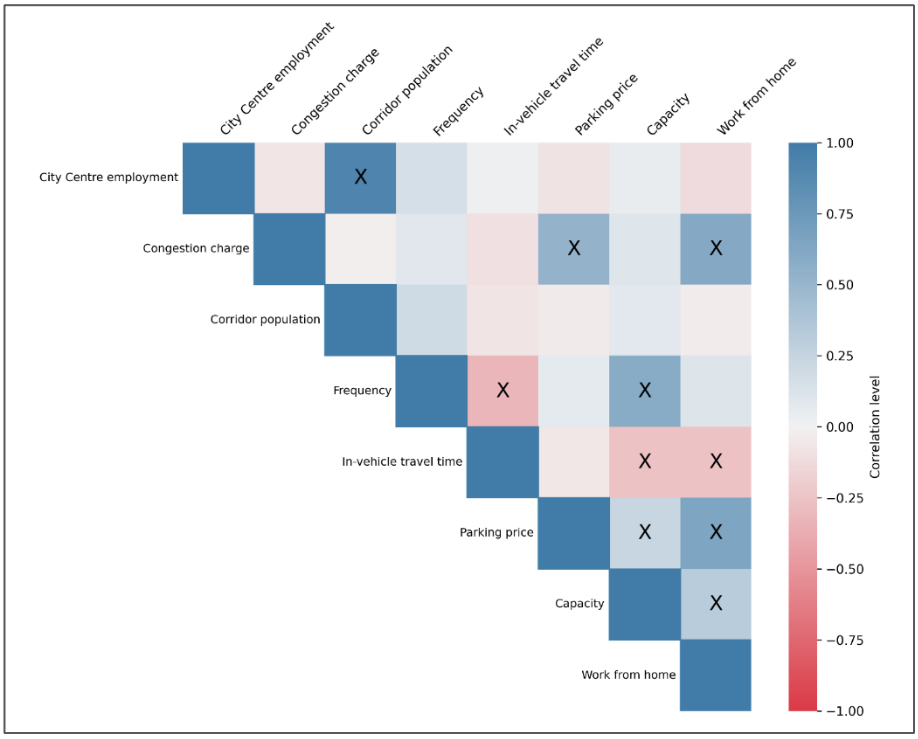

Notes: "X" denotes correlations that were statistically significant at the 1% level.

The multicollinearity evident in Figure 3 arises because model runs were initially designed to meet policy needs from the business case, rather than to optimise the estimation of regression models like we do here. For example, the 0.52 correlation between parking pricing and congestion charges does not imply a causal relationship but instead reflects how the scenarios were designed. These correlations serve to reduce the precision of estimated parameters, making it more difficult for us to isolate the individual impacts of some important policy variables, such as congestion pricing.

Transport model runs could be undertaken in a more scientific way. Orthogonal fractional factorial designs – like those used in stated choice experiments (Louviere, et al., 2000; Groemping & Morgan-Wall, 2025) – could be used to systematically vary the inputs into FSMs in a way that was designed to increase variance and reduce multicollinearity while using a small number of runs. Our ongoing research uses a fractional factorial design to identify 36 model runs, which in turn generate sufficient variance to support regression modelling for this case study.

Second, the illustrative case presented in this paper required us to make assumptions on the statistical distribution of inputs, such as population and employment, and associated distributional parameters. Although our assumptions were informed by our professional understanding of the nature of the uncertainty that characterises these outcomes, such assumptions would ideally be grounded in empirical analyses. Using empirical data to characterise uncertainty in inputs would help to strengthen confidence in the results of the Monte Carlo simulations in Step 2, which here rely primarily on experience and engagement and are, by definition, ad-hoc. The latter is something this approach seeks to avoid and will be a focus of our own future research.

Third, and on a related point, we currently model uncertainty in inputs by repeatedly drawing independent samples from the assumed underlying statistical distributions. In practice, we might want to allow for correlations between the underlying inputs, such as population and employment, which could then be accounted for directly in the sampling process. Formally allowing for such correlations where they are supported by empirical evidence would further underscore the added value of our proposed approach compared to more conventional treatments of uncertainty, such as sensitivity testing, in which correlated effects are rarely considered in a systematic way.

CONCLUSIONS

Transport planning and investment processes usually rely on travel demand forecasts from transport models. Although these models have been continuously refined through decades of research and real-world application, they usually generate deterministic ("point") estimates that do not provide insights into uncertainty. As forecasts often extend decades into the future, we suggest that uncertainty is an inherent aspect of travel demand forecasts that warrants more attention.

Directly using transport models to quantify uncertainty, however, is relatively impractical. In large urban areas with complex transport networks and travel patterns, individual model runs can take several days to converge. As a result, using transport models to directly quantify uncertainty, for example, by repeatedly varying inputs and re-running the model to generate a distribution of travel demand forecasts currently strikes us as being somewhat impractical.

In response to these observations, we have proposed an approach to quantifying input uncertainty that combines Bayesian regression models with Monte Carlo methods. This approach allows us to use travel demand models in a parsimonious way to quantify the uncertainty in travel demand forecasts associated with selected inputs. We illustrate the approach by applying it to a case study, where we find uncertainty in a limited subset of inputs causes estimated demands to vary by approximately 5-15%. This variation is significant and highlights, in our view, the merits of incorporating probabilistic insights into transport planning and investment processes. We also note several limitations with our current approach and highlight opportunities for further research. These include adopting a structured ("orthogonal") approach to the design of model runs, using empirical analyses to inform statistical assumptions, and allowing for correlations between inputs.

Despite these opportunities for improvement, we nonetheless consider the proposed approach to be sufficiently well developed to be applied in practice. Using these methods to systematically account for uncertainty in travel demand forecasts has the potential to inform transport planning and investment decisions. Uncertainty is, in our view, inherent to planning: It should not be managed via ad hoc ex post sensitivity tests but rather explicitly considered as a fundamental upfront design consideration for all major transport interventions.

REFERENCES

Auckland Light Rail, 2024. Draft-Corridor-Business-Case-Economic-Case. [Online].Available at: https://www.transport.govt.nz/assets/Uploads/Redacted-2.-Draft-Corridor-Business-Case-Economic-Case-marked-up-for-redaction-3.pdf

Balcombe, R. et al., 2004. The demand for public transport: A practical guide, United Kingdom: TRL Limited.

Groemping, U. & Morgan-Wall, T., 2025. CRAN Task View: Design of Experiments (DoE) & Analysis of Experimental Data. [Online]

McNally, M. G. (2007). The four-step model. In Handbook of transport modelling (Vol. 1, pp. 35-53). Emerald Group Publishing Limited.

Lieswyn, J. (2011). Probabilistic risk analysis in transport project economic evaluation (Master of Engineering thesis, University of Canterbury). UC Research Repository.

Louviere, J., Hensher, D. & Swait, J., 2000. *Stated Choice Methods: Analysis and Applications.* s.l.:Cambridge University Press.

Wallis, I., 2004. Review of Passenger Transport Demand Elasticities, Wellington New Zealand: Transfund New Zealand Research Report No. 248.

Willumsen, L., 2014. Better Traffic and Revenue Forecasting. London: Maida Vale Press.

Willumsen, L. G. & Ortúzar, J. d. D., 2016. Transport planning. In: Handbook on Transport and Urban Planning in the Developed World. s.l.:Edward Elgar Publishing Limited.

AUTHOR CONTRIBUTION STATEMENT

Dr Stuart Donovan was responsible for modelling, analyses, and communication of results. Peter Clark was responsible for conceptualisation, data collection, preliminary modelling, and communicating of results.