

Assessing Uncertainty About the Benefits of Transportation
Infrastructure Projects Using Bayesian Melding: Application to
Seattle's Alaskan Way Viaduct

Hana Ševčíková, Adrian E. Raftery and Paul A. Waddell
University of Washington

Working Paper no. 90
Center for Statistics and the Social Sciences
University of Washington

March 26, 2009

Abstract

Uncertainty is inherent in major infrastructure projects, yet standard practice ignores it. We investigate the uncertainty about the future effects of tearing down the Alaskan Way Viaduct in downtown Seattle, using an integrated housing, jobs, land use and transportation model, on outcomes including average commute times. Our methodology combines the urban simulation model UrbanSim with the travel model Emme/3. We assess uncertainty using Bayesian melding, yielding a full predictive distribution of average commute times on 22 different routes in 2020. Of these routes, 14 do not include the viaduct and eight do. For the 14 base routes that do not include the viaduct, the predictive distributions overlap substantially, and so there is no indication that removing the viaduct would increase commute times for these routes. For each of the eight routes that do include the viaduct, the 95% predictive interval for the difference in average travel times between the two scenarios includes zero, so there is not strong statistical support for the conclusion that removing the viaduct would lead to any increase in travel times. However, the median predicted increase is positive for each of these routes, with an average of 6 minutes, suggesting that there may be some measurable increase in travel time for drivers that use the viaduct as a core component of their commute.

Contents

1	Introduction	1
2	UrbanSim with Integrated Travel Model	2
2.1	Land Use Models	2
2.2	Travel Model	4
2.3	Integrated Land Use and Travel Model	4
3	Policy Question: Seattle’s Alaskan Way Viaduct Replacement	4
4	Bayesian Melding Method and its Application	6
4.1	Bayesian Melding Method	6
4.2	Results from Prior Research	7
4.3	Data	8
4.4	Prior, Likelihood and Posterior distribution of the Land Use Model	8
4.5	Calibration of the Travel model	10
5	Results from Integrated Land Use and Travel Model	10
6	Discussion	13
7	Acknowledgments	17

List of Tables

1	Estimates for bias and variance, respectively, obtained from the run based on the prior mean of Θ	10
---	--	----

List of Figures

1	Integrated use of UrbanSim and the Travel Model.	4
2	Seattle’s Alaskan Way Viaduct.	5
3	Illustration of the Bayesian melding method for deterministic models	7
4	Map of seven commuter routes that do not include the Alaskan Way Viaduct	9
5	Calibration of the travel times	11
6	Nested simulation scheme for simulating the posterior distribution of average travel time in 2020 by Bayesian melding.	12
7	Map of routes that include the Alaskan Way Viaduct	13
8	Posterior predictive distributions of average travel times in 2020 for the routes that do not include the viaduct	14
9	Posterior predictive distributions of average travel times in 2020 for the routes that include the viaduct	15
10	Posterior predictive distributions of differences between average travel times for the two scenarios	16

1 Introduction

Major infrastructure investments routinely cost billions of dollars to construct, and are rife with political conflict over the costs, the benefits, and their distribution over often competing stakeholders. Transportation projects such as highways, bridges, and light rail systems are lightning rods for political controversy. One would imagine that with the advanced state of the art in modeling travel behavior and traffic flow dynamics, these questions would be relatively straightforward to address using metrics to measure the benefit of the project in terms of its effects on travel times, for example, as compared to the costs of the project. The state of the practice is substantially less informed than this, unfortunately.

We use a motivating case study that will be a central part of this paper to illustrate. The Alaskan Way Viaduct, an elevated highway constructed in the 1950's along the downtown Seattle waterfront, is often compared to a similarly designed elevated Embarcadero Freeway along the waterfront of downtown San Francisco that was eventually demolished in 1991 after being damaged by the Loma Prieta earthquake in 1989. The Alaskan Way Viaduct was damaged by the 2001 Nisqually earthquake and has been for the past several years a point of controversy among government officials ranging from the Mayor of Seattle to the Governor of Washington, about how to eliminate the risk of catastrophe from a collapse of the elevated highway in the next earthquake, and how to replace the facility in a way that appeases competing interests. It will be a costly project, with estimates ranging from \$2.5 billion to well over \$4 billion for various options. In January 2009, the Mayor, King County Executive and Governor jointly announced an agreement to replace the viaduct with a tunnel option that would cost approximately \$4.25 billion.

But whether we need as much, or more, transportation capacity as is now carried by the viaduct, has been a point of contention with environmentalists and transit advocates, and may be at odds with other objectives adopted in state policy, such as achieving a dramatic reduction in greenhouse gas emissions over the next several decades. Other advocates claim that if we replace the viaduct with any alternative that has less capacity to move high volumes of traffic past the downtown area, as an alternative to the I-5 corridor, the highway system will be slowed to a crawl, with massive costs in delays to commuters and to businesses that depend on this access.

No one questions whether the viaduct should be removed and replaced with some kind of facility. The debate is over what the replacement should be, and how much capacity it should carry. At the heart of this debate, then, are quantities that are uncertain: the benefits in terms of travel time, and the costs of alternative projects. While there is a reasonable basis for assessing uncertainty about costs, there is relatively little guidance on how to assess the uncertainty about benefits. In fact, the models used to predict travel patterns based on alternative transportation networks, are designed to run to a deterministic equilibrium, and traffic modelers rely heavily on the certainty that repeated runs of the model produce consistent results — in other words, that they contain no uncertainty (Boyce 1984).

In fact, these decisions depend on information that is uncertain, and the public has a direct interest in better understanding how likely it is that spending several billion dollars will actually solve a problem they care about. Our paper contributes to this debate and provides a foundation for further work to incorporate uncertainty more systematically into the planning process, and into public deliberation about large, expensive projects with long-term impacts. There is limited prior work that examines the issue of uncertainty in the context of integrated transportation and land use models (Kockleman 2002; 2003, Clay and Johnston 2006, Ševčíková et al. 2007), and in network capacity and design (Sumalee et al. 2009). Our paper extends the literature in this area by using a principled statistical method to calibrate uncertainty in an integrated land use and transportation model system, and hence to assess the uncertainty of specific metrics that reflect

the potential benefits of a major transportation facility. To our knowledge, this is the first research to use Bayesian melding to assess the uncertainty about the travel time impacts of alternative investments in major transportation facilities.

The debate centers on the question of whether it is possible to reduce capacity by removing a waterfront highway such as the Alaskan Way Viaduct, without greatly increasing travel times for commuters and commercial vehicles. On its face this seems unlikely to be possible, but some of the literature that addresses induced demand from capacity expansion, such as Downs (2004), suggests an argument that it may be. Downs coined the term ‘triple-convergence’ to describe the propensity for commuters to take advantage of increases in roadway capacity and temporarily faster speeds by changing routes, times of travel and modes of travel, in order to take advantage of the relative increase in speed of travel on the improved highway at peak hour by single-occupancy vehicles. The question of induced demand has rarely been raised in the context of a capacity reduction, but there is nothing inherent in the reasoning that would prevent it from applying to such a case. In the event of a capacity reduction, such as the loss of a highway, travelers would presumably make short-term choices that would shift away from the relatively higher cost route, time and mode to those that become relatively less expensive. Consider this a case of ‘reduced demand’.

In the longer term, of course, persons, households and businesses can adapt to changes in accessibility by changing their locations. These longer-term induced demand or reduced demand effects may be at least as big as the short term effects described by triple-convergence (Downs 2004). In this paper, we set out to explore these questions using a land use model, UrbanSim (Waddell 2002, Waddell et al. 2003; 2007), coupled with a four-step transportation model implemented by the Puget Sound Regional Council (PSRC). The approach we develop is Bayesian melding, a methodology initially developed to calibrate uncertainty in deterministic model systems by Raftery et al. (1992; 1995) and Poole and Raftery (2000), and recently adapted to stochastic models by Ševčíková et al. (2007).

Our contribution is to harness the Bayesian melding approach to calibrate uncertainty in a combined land use and transportation model system, and to use the calibrated system to make inferences about the effects on travel times of two different alternatives of the Alaskan Way Viaduct. As the objective of this paper is not to make a definitive assessment of the specifics of the viaduct project design, we approximate the alternatives by modeling one as having the same capacity as the existing viaduct, and the other as a worst-case scenario in which the viaduct is simply removed, and no mitigation is done in terms of local street configuration and operations or of transit service in this corridor. The intent is to demonstrate on a real-world, and still timely case, the use of uncertainty analysis to inform the policy debate such as this. It should be broadly applicable as a methodology to a much wider set of problems.

The paper proceeds with a brief description of the models, since their internal construction is not the focus of this paper, and details of the models used in the analysis are available in the provided citations. We then present the Bayesian melding approach developed for application to this case study, and close with a discussion of the results and implications for further research.

2 UrbanSim with Integrated Travel Model

2.1 Land Use Models

UrbanSim is an urban simulation model operational in several urban areas in the United States (Waddell 2002, Waddell et al. 2003; 2007). The system is implemented as a set of interacting models that represent the major actors and choices in the urban system, including households choosing residential locations, business choices of employment location, worker choices of jobs and developer

choices of locations and types of real estate development. The model system microsimulates the annual evolution in locations of individual households and jobs, including the connection between them, and the evolution of the real estate within each individual geography as the result of actions by real estate developers.

Our application of UrbanSim operates on parcel level. It is configured to run the following models:

1. **Real estate price model:** predicts prices of parcels, using a hedonic regression model.
2. **Expected sale price model:** predicts prices of possible real estate proposals, using a hedonic regression model.
3. **Development proposal choice model:** chooses real estate proposals to be built (including redevelopment), using weighted random sampling based on a predicted return on investment (ROI).
4. **Building construction model:** demolishes buildings (for redevelopment) and builds new buildings according to the chosen proposals.
5. **Household transition model:** creates and removes households and updates the set of persons accordingly. It is based on random sampling and is driven by macroeconomic predictions.
6. **Employment transition model:** creates and removes jobs, using random sampling, and is driven by macroeconomic predictions.
7. **Household relocation choice model:** determines households for moving, using a logit model.
8. **Household location choice model:** locates moving households into buildings, using a multinomial logit model.
9. **Employment relocation model:** determines jobs for moving using weighted random sampling.
10. **Employment location choice model:** locates moving jobs into buildings, using a multinomial logit model.
11. **Work at home choice model:** simulates workers decision to work at home or out of home. It is based on a logit model.
12. **Workplace relocation choice model:** simulates workers decision to change job. It is based on a logit model.
13. **Workplace choice model:** assigns jobs to workers, using a multinomial logit model with sampling alternatives.

Several of the models require coefficients which are obtain by estimating using observed data and Maximum Likelihood Estimation (MLE) of multinomial logit models based on a Random Utility Maximization framework (McFadden 1974; 1978; 1981). Most of the models are stochastic, and involve Monte Carlo sampling of choice outcomes conditional on a probability generated from a Multinomial Logit Model (MNL). A simulation starts to operate on observed data (so called base year data) about households, persons, jobs, buildings, parcels, zones etc. Each iteration of the model system modifies the data and is considered as a prediction for the particular year.

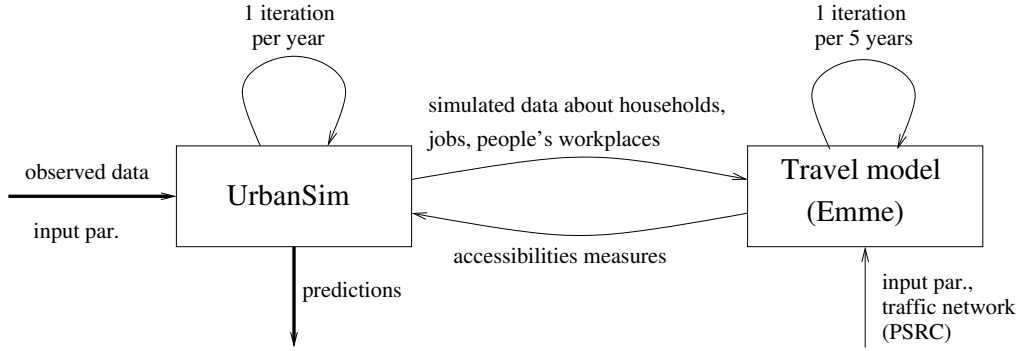


Figure 1: Integrated use of UrbanSim and the Travel Model.

2.2 Travel Model

The travel model used by the PSRC is a state of the practice four-step travel model, in the early stages of transition to a full activity-based travel model system. So far, only the trip generation step has been replaced by an activity generation model, and the rest of the model system operates as a traditional aggregate travel model with destination choice implemented as a gravity model, mode choice as a multinomial logit model, and assignment using Emme/3, with feedback of congested travel times to mode choice. Full documentation of the base model system is available from the PSRC web site (www.psrc.org).

The travel model runs on a zonal system that contains approximately 1,000 zones to cover the Central Puget Sound, consisting of King, Kitsap, Pierce and Snohomish Counties. It is implemented in the Emme/3 software platform, and requires approximately 16 hours to complete one simulation year, including iteration for convergence on assignment, on a high-end desktop computer.

2.3 Integrated Land Use and Travel Model

The integration of UrbanSim and the Travel Model can be explained using Figure 1. UrbanSim is run on an annual basis, i.e. one iteration of the full set of models simulates a land use evolution in one year. Due to the very high run times of the travel model, we run Emme/3 only once in five years of UrbanSim runs. The travel model uses the output from UrbanSim about households, jobs and people’s workplaces. In addition, it has its own input parameters and it operates on a network provided by the PSRC. In turn, several UrbanSim models use accessibilities measures computed by the travel model, such as travel times or log sums. They are used mainly as predictive variables in the Employment location choice model.

3 Policy Question: Seattle’s Alaskan Way Viaduct Replacement

The Alaskan Way Viaduct, built in 1953, is an elevated section of State Route 99 that runs along the Elliott Bay waterfront in Seattle’s Industrial District and downtown Seattle (see Figure 2). In Figure 4, the viaduct is shown as a black solid line. The viaduct was damaged in the 2001 Nisqually earthquake and since then continuing settlement damage has been discovered (WSDOT 2005). In 2002, the Washington State Department of Transportation (WSDOT) together with the City of Seattle, the Federal Highway Administration and King County have launched a program that would lead to a replacement of the viaduct (WSDOT 2004). Since then, many replacement




Figure 2: Seattle's Alaskan Way Viaduct.

concepts and designs have been evaluated, and these were narrowed down in 2008 into three hybrid solutions:

- Surface and transit option: the viaduct is removed; significant improvements in surface and transit conditions.
- Elevated structure: the viaduct is rebuilt, but with current design standards which would require a larger structure.
- Tunnel option: a four-lane 2-mile underground tunnel with improvements to the seawall and other streets.

One of the main objections raised by critics of the surface transit option has been a fear that it would produce traffic jams and drastically increased travel times on routes along the viaduct, as well as on I-5, which runs parallel to the viaduct on the east side of downtown. Though proponents of the surface transit option have pointed out that the demolition of the Embarcadero Freeway in San Francisco did not cause significant traffic problems, the viaduct carries considerably more traffic. The viaduct carries approximately 110,000 cars per day, whereas the Embarcadero Freeway carried around 70,000 cars per day before its demolition. Further, the geography of Seattle, constrained by water on its east and west sides, means that the I-5 corridor is the only major north-south freeway through Seattle. It is thus legitimate to ask whether reducing capacity on the viaduct would make the already bad I-5 traffic much worse.

WSDOT released a study that compared various transportation measures for eight different scenarios, see WSDOT (2008) and WSDOT (2007) (meetings from November 13 and November 24

2008). These measures included traffic volume, pattern and modes of travel as well as the quality of those trips as measured by forecast travel times during various periods of the day. The baseline was set to the year 2015 and the study area was limited to the city center of Seattle. The land use data used as inputs for the travel model incorporated a growth in the downtown area. Their key findings in terms of travel time were that a trip through the city from the north to the south side at an AM peak would be approximately 10 minutes longer if there is the surface option implemented as opposed to an elevated structure. Their model did not take into account changes in land use over time, including changes in real estate prices. Moreover, it provides point predictions without any sense about the uncertainty of the results.

In January 2009, the Washington State Governor together with the Mayor of Seattle announced an agreement to pursue the tunnel option. Though a decision on the alternative appears to have been concluded (there have been several changes in the decision process along the way, so this may or may not be the final outcome), we think this case is still relevant for our study, as it highlights the issue of uncertain benefits from large-scale infrastructure projects.

We are interested in comparing changes in travel times over time resulting from different viaduct replacement options. Since our motivation is the development of a better method to inform such decisions, we are less interested in the fine points of the design of the alternatives. Rather, we develop two alternatives that should provide a suitable bracket for the alternatives that have been considered. For simplicity, we use as a base alternative a network that matches the current capacity of the existing viaduct. Whether it comes in the form of a tunnel or a replacement elevated structure is not material to this analysis.

For the other alternative, we take a worst-case scenario that should be dramatically worse in terms of effects on travel times than the surface transit option that has been under consideration: for this worst-case scenario, we simply remove the viaduct in 2010, and provide no mitigation in terms of improved transit service, or improvements to local streets in downtown. It is truly an unrealistically worst case. The rationale for this is that we want to examine whether there is a large enough difference in travel times between these two cases to offset the uncertainty in the analysis of the travel time benefits. One would like to think that the results generate confidence that the investment of more than \$4 billion would improve travel times, over the alternative that was used in the Embarcadero case: simply removing the elevated highway and connecting downtown to the waterfront.

Thus, to summarize, we compare two scenarios:

1. **Capacity-Neutral Replacement** which is our baseline. We use the travel model networks provided by PSRC for years 2005, 2010, 2015 and 2020. In terms of travel times, this scenario approximates a situation in which either the viaduct is rebuilt or a tunnel is built.
2. **Worst-Case: Demolish Viaduct** in 2010. Here we remove links from the 2010, 2015 and 2020 networks that represent the viaduct.

4 Bayesian Melding Method and its Application

4.1 Bayesian Melding Method

Bayesian melding was proposed by Raftery et al. (1992; 1995) and Poole and Raftery (2000) as a way of putting the analysis of deterministic simulation models on a solid statistical basis. The method was modified and applied to stochastic models by Ševčíková et al. (2007), specifically to urban simulation models.

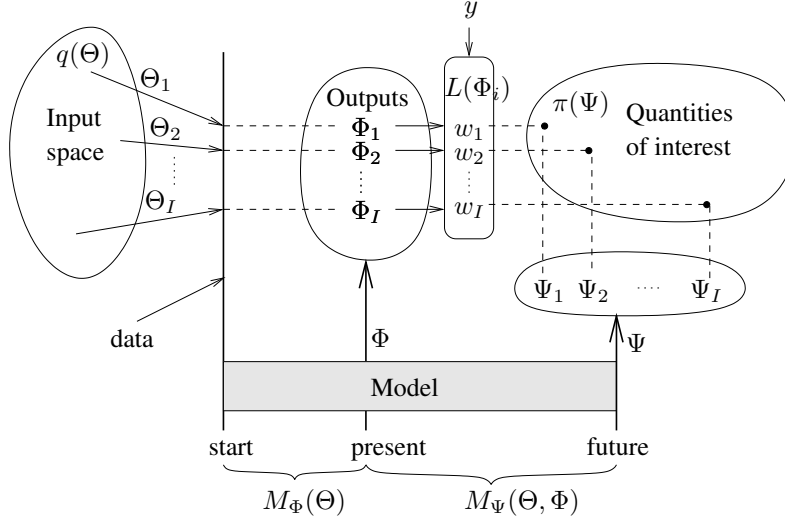


Figure 3: Illustration of the Bayesian melding method for deterministic models. The uncertain model inputs, Θ , refer to the starting time of the simulation, and the outputs, Φ and the data relevant to the outputs, y , are observed at the “present” time, while the quantities of interest, Ψ , refer to the future. The quantities Θ_i , Φ_i and Ψ_i refer to the i -th simulated values of inputs, outputs and quantities of interest, respectively.

A simple version of the original method for deterministic models is summarized in Figure 3. There is a prior distribution of model inputs $q(\Theta)$ from which we draw input values Θ_i for $i = 1, \dots, I$. The model runs I times from the starting point to the present and for each input Θ_i it produces as output the quantity of interest, Φ_i . The model can be viewed as a mapping, M , from the space of inputs to the space of outputs, which we denote by $\Phi = M_\Phi(\Theta)$. The “present” time is defined as a time point for which we have observed data available. We use the observed data, denoted by y , to compute a weight w_i for each input Θ_i : $w_i = L(\Phi_i)$. Here, $L(\Phi_i)$ is the likelihood of the model outputs given the observed data, $L(\Phi_i) = \text{Prob}(y|\Phi_i)$. For each of the I runs, the model is run forward until a future time for which we make a prediction. The results of the i th model run are denoted by Ψ_i . The posterior distribution of Ψ is approximated by a discrete distribution with values Ψ_i having probabilities proportional to w_i .

The method was extended to stochastic mechanistic models such as UrbanSim by Ševčíková et al. (2007). The main change was that the conditional distribution of the model outputs Φ given the model inputs Θ , which is a point mass at $M_\Phi(\Theta)$ for deterministic models, became a probability distribution. This distribution had two components, one reflecting the stochastic nature of the model outputs, and the second reflecting model error. Details can be found in Ševčíková et al. (2007).

4.2 Results from Prior Research

In Ševčíková et al. (2007) we applied the method to an UrbanSim application for the Eugene, Oregon region. We were able to determine the posterior predictive distribution of the numbers of households in each of the 295 traffic analysis zones and in any aggregation of those¹. Our starting

¹In that research, the travel model was not included in the analysis, whereas this current research incorporates both the land use and transportation model components, and extends the method to the assessment of uncertainty

point was the year 1980, the “present” time was 1994 and the “future” time was 2000. Using observed data in 2000 we were able to validate the results. Our main conclusions were:

- The Bayesian melding approach provided well calibrated results, while simple multiple runs (reflecting the stochastic variation in model output, but not model error) underestimated uncertainty.
- A square root transformation of the quantity of interest (number of households) yielded an approximately constant variance of the model error.
- Variation of the model inputs and random seed did not account for much of the uncertainty.

4.3 Data

Our simulation region is the Puget Sound area, and our starting point, or base year, is 2000. We have detailed information about the Puget Sound area 2000, which includes 1, 282, 940 households, 1, 608, 426 workers, 1, 849, 447 jobs, 1, 008, 869 buildings, 1, 177, 140 parcels, and 938 traffic analysis zones (TAZ).

We also have less detailed data observed in 2005, taken as the “present” time. This includes the numbers of households in each TAZ and the numbers of jobs in each TAZ divided into 8 groups: mining; construction; manufacturing; wholesale trade, transportation, and utilities (wctu); retail trade; financial, professional, health and other services (fires); education; and government. These will be our calibration data y for the land use model.

In order to calibrate the travel model output, we obtained observed travel times for selected routes in 2005 from the Washington State Department of Transportation (<http://depts.washington.edu/hov>). These are annual averages over weekdays in five-minute periods, which we averaged over the AM peak (6:00am – 9:00am) in order to do a direct comparison with the travel model outputs. We chose 7 non-overlapping popular commuter routes, i.e. 14 trips, for which average travel times were available; see Figure 4.

4.4 Prior, Likelihood and Posterior distribution of the Land Use Model

We first extend the statistical model of Ševčíková et al. (2007) on which the likelihood function $L(\Phi_i) = \text{Prob}(y|\Phi_i)$ is based, for use with multiple quantities of interest, as follows:

$$(y_{kl}|\Theta = \Theta_i) = \mu_{ikl} + a_l + \epsilon_{ikl}, \text{ where } \epsilon_{ikl} \stackrel{iid}{\sim} N(0, \sigma_{il}^2), \quad (1)$$

for $i = 1, \dots, I$, $k = 1, \dots, K$ and $l = 1, \dots, L$. Here i indexes the simulation run, k indexes the zone, and the index l refers to the l -th quantity of interest. The quantity μ_{ikl} is the expected value of y_{kl} under the model given Θ_i , ϵ_{ikl} denotes the model error, and a_l is the overall bias in the model predictions of the l th output. The variance σ_{il}^2 and bias a_l are estimated by their sample equivalents: $\hat{\sigma}_{il}^2 = \frac{1}{K} \sum_k (y_{kl} - \hat{a}_l - \hat{\mu}_{ikl})^2$, and $\hat{a}_l = \frac{1}{IK} \sum_{i,l} (y_{kl} - \hat{\mu}_{ikl})$, where $\hat{\mu}_{ikl}$ is the predicted value of y_{kl} from the i th simulation run.

This yields a conditional predictive distribution of our quantity of interest:

$$y_{kl}|\Theta_i \sim N(\hat{a}_l + \hat{\mu}_{ikl}, \hat{\sigma}_{il}^2). \quad (2)$$

in evaluation of infrastructure alternatives.




Figure 4: Seven commuter routes none of which includes the Alaskan Way Viaduct. The viaduct is shown in black.

We then have

$$w_i \propto p(y|\Theta_i) = \prod_{l=1}^L \prod_{k=1}^K \frac{1}{\sqrt{2\pi\hat{\sigma}_{il}^2}} \exp \left[-\frac{1/2(y_{kl} - \hat{a}_l - \hat{\mu}_{ikl})^2}{\hat{\sigma}_{il}^2} \right]. \quad (3)$$

The quantities $\hat{\sigma}_{il}^2$ and \hat{a}_l are estimated at the “present” time $t_1 = 2005$. The marginal distribution of the l -th quantity of interest, Ψ_{kl} , in the year $t_2 = 2020$, is given by a mixture of normal distributions, as follows:

$$\pi(\Psi_{kl}) = \sum_{i=1}^I w_i N(\hat{a}_l b_a + \Psi_{ikl}, \hat{\sigma}_{il}^2 b_v), \quad k = 1, \dots, K, l = 1, \dots, L. \quad (4)$$

Here, b_a and b_v denote propagation factors of the bias and the variance over the time period $[t_1, t_2]$.

In this application, the long runtime of the travel model made it infeasible to do a large number of runs of the travel model. Also, we found that the results of UrbanSim for numbers of households and jobs were relatively insensitive to the values of Θ_i drawn from the prior (results not shown). Thus the contribution of uncertainty about the UrbanSim inputs Θ to overall uncertainty about average travel times was small. In particular, the variation in σ_{il}^2 between runs was small, and so we used a single estimate, $\hat{\sigma}^2$, using the run based on the prior mean of Θ , estimated from external data. Results (computed on the square root scale) are shown in Table 1.

In addition, we were interested in comparisons between scenarios, and assuming that the propagation factors were the same for both scenarios allowed us to ignore them and set them both equal

l	measure	\hat{a}_l	$\hat{\sigma}_{1l}^2$
1	households	-0.02	7.2
2	mining	-0.21	4.9
3	construction	0.25	20.2
4	manufacture	-0.80	15.4
5	wtcu	-0.08	24.8
6	retail	0.07	21.0
7	fires	0.38	35.4
8	education	-0.57	28.7
9	government	0.39	37.5

Table 1: Estimates for bias and variance, respectively, obtained from the run based on the prior mean of Θ .

to 1. Together, these considerations allowed us to approximate (4) by the simpler equation

$$\pi(\Psi_{kl}) = \frac{1}{I} \sum_{i=1}^I N(\hat{a}_l + \Psi_{ikl}, \hat{\sigma}^2), \quad k = 1, \dots, K, \quad l = 1, \dots, L. \quad (5)$$

For priors, we used the same approach as Ševčíková et al. (2007). For input parameters that were estimated by multinomial logistic regression or by hedonic regression from external data, we used the multivariate normal distribution $MVN(\hat{\Theta}, SE(\hat{\Theta})^2)$, with mean $\hat{\Theta}$, the estimator of Θ , and with as variance matrix the diagonal matrix with diagonal entries equal to the squares of the standard errors of the parameters. For mobility rates used in the Employment relocation model, we used the normal distribution $N(\hat{r}, (\frac{\hat{r}(1-\hat{r})^2}{n}))$, truncated at zero, where \hat{r} is an estimate of the rate r and n is the number of observations from which \hat{r} was obtained.

The land use model uses regional control totals for number of households and jobs obtained from external sources. We kept the control totals constant, and so the results are conditioned on these totals.

4.5 Calibration of the Travel model

Due to the complexity of the input parameters and the long run-times of the travel model, we assessed uncertainty about the travel model by a simple calibration procedure. In Figure 5, we plotted the simulated average travel times for the different commutes in $t_1 = 2005$ against the observed average travel times (obtained as described in Section 4.3). As can be seen, the travel model overestimates the travel times.

We found that, given the simulated average travel time T_{sim} , the conditional distribution of the observed average travel time, T , was well represented by a normal distribution on the logarithmic scale with an additive bias:

$$\log(T) \sim N(\log(T_{sim}) - 0.70, 0.14^2). \quad (6)$$

5 Results from Integrated Land Use and Travel Model

The posterior distribution of the resulting travel time T is given by

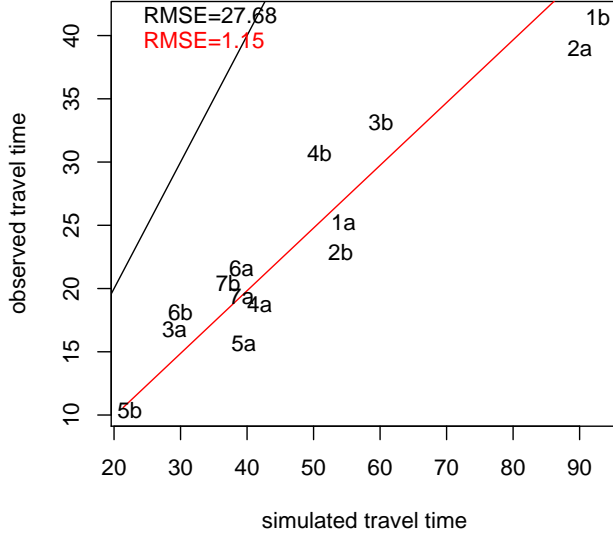


Figure 5: Calibration of the travel times. The solid line is the $y = x$ diagonal; the dashed line is $y = e^{-0.70}x$, corresponding to the calibration in equation (6). The points are numbered according to the routes in Figure 4.

$$\pi(T) = \int p(T|\Psi, \Theta) p(\Psi|\Theta) p(\Theta) d\Theta d\Psi, \quad (7)$$

where $p(T|\Psi, \Theta)$ is given by (6) and $p(\Psi|\Theta)$ is simulated from by running UrbanSim with inputs Θ and applying equation (5). Because we are interested in comparisons between scenarios, we ignore the propagation factors in (4), which we assume will be the same for both scenarios.

For both the baseline (viaduct) scenario and the no-viaduct scenario, we evaluated the integral in (7) by simulation, using the nested simulation scheme depicted in Figure 6. Given the long time needed to run the travel model, we approximated the integral over Θ by simulating a small number, I , of values of Θ from its prior distribution, and approximating the integral by an equally-weighted discrete distribution over $(I + 1)$ values of Θ , namely the I simulated values and the point estimate from external data, as in equation (5). As discussed above, this may overestimate this source of uncertainty, since it does not allow for the additional information about Θ from the 2005 data, but the estimated uncertainty from this source was small in any event, and so we found this approximation adequate.

To simulate a value of the outputs Ψ given a value of Θ , we ran UrbanSim for the first five years of the simulation period (2000–2005), and then we ran the travel model. This was repeated for each five-year period until 2020.

For each simulated value of Θ , we simulated J values of the set of outputs Ψ (numbers of households and jobs for each zone in 2020) by the method described in Section 4. Finally, for each simulated value of Ψ , we simulated N values of T from (6). We used $I = 3$, $J = 5$ and $N = 1,000$. Note that $\hat{\sigma}^2$ was obtained only once, from the run with the prior mean for Θ estimated from the

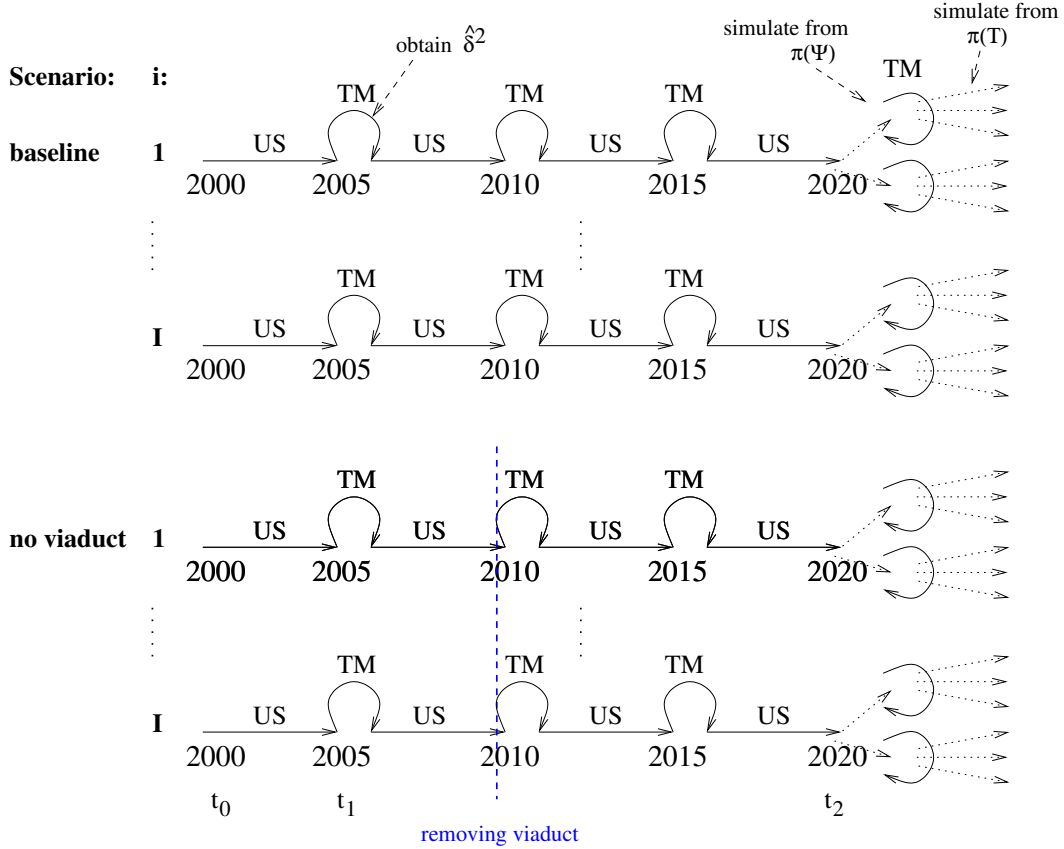


Figure 6: Nested simulation scheme for simulating the posterior distribution of average travel time in 2020 by Bayesian melding.

external data, and reused in all remaining runs.

Results are shown in Figures 8 and 9. The figures show the posterior distributions of average travel times for the two scenarios: baseline in grey, no-viaduct in red. Figure 8 contains the seven routes (in both directions) from Figure 4 that do not contain the viaduct. Figure 9 contains eight additional routes that go (or would go) directly through the viaduct as shown in Figure 7.

From Figure 8 it is clear that the posterior predictive distributions of average travel times under the two scenarios overlap substantially, so that our analysis does not indicate that removing the viaduct would have any effect on average travel times for commuter routes that do not include the viaduct. For the routes that do include the viaduct, Figure 9 shows that the posterior distributions still overlap, but not completely.

To investigate further whether our results indicate an effect of removing the viaduct on average travel times, we calculated the posterior predictive distribution of the *difference* between average travel times under the two scenarios in 2020. These are shown for all routes considered in Figure 10. The seven base commutes that do not include the viaduct are in the upper part of the figure, and it is again clear that our analysis does not indicate any effect of removing the viaduct for these routes, since zero is close to the center of all the distributions.

For the routes that do include the viaduct the situation is less clear. The 95% predictive intervals for all of these routes includes zero, so our simulation results do not clearly indicate an effect of removing the viaduct. On the other hand, the median change for all eight routes that

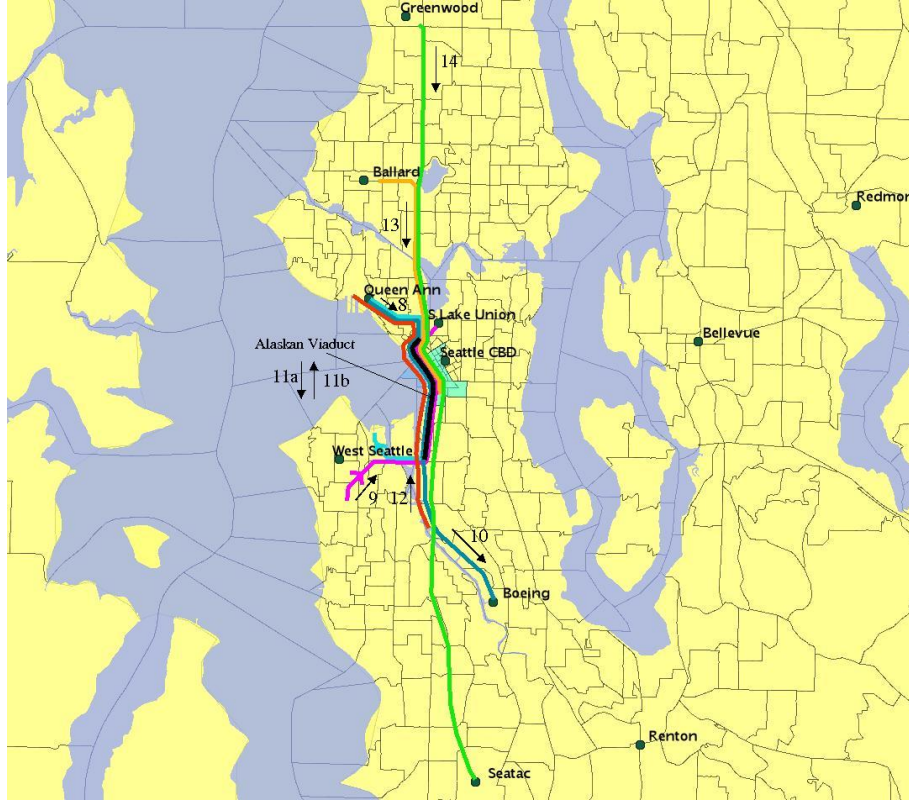


Figure 7: Routes that include the Alaskan Way Viaduct. Route 11 (shown in black) is the viaduct itself: 11a goes from the north end to the south end of the viaduct, while 11b goes from south to north.

contain the viaduct is positive, ranging from 1.5 to 9.2 minutes, and averaging 6.1 minutes. The median predicted change for traveling the viaduct alone from north to south (route 11a) is 5.7 minutes, but the predictive interval contains zero.²

6 Discussion

In the Seattle Times of November 14, 2008 an article of Gilmore (2008) reported on the WSDOT (2008) study. It indicated that if the viaduct were replaced by another elevated highway in 2015, drivers going from Greenwood in North Seattle to SeaTac International Airport (our route 14) would arrive 10 minutes sooner than if the replacement were a surface boulevard. Furthermore, from Ballard to south of Downtown (our route 13) would be 13 minutes faster and drivers on our route 12 would save 10 minutes on an elevated highway. These point estimates ignore any uncertainty involved in the models used to generate them, and thus could mislead the public into having an unwarranted degree of confidence in the benefits of making these investments. The point estimates for routes 12 and 14 fall into our prediction interval (Figure 10), whereas the 13 minutes

²Note that the routes that contain the viaduct all overlap, in some cases substantially, and so the posterior predictive distributions for different routes are not independent. As a result, it is not possible to view these eight posterior predictive distributions as independent samples from a distribution and carry out a standard statistical test on the average value.




Figure 8: Posterior predictive distributions of average travel times in 2020 for each of the seven routes from Figure 4 in both directions. The posterior distributions are represented by boxplots, with the box containing the interquartile range with the median marked in the middle, and the whiskers covering the 95% posterior confidence interval. The baseline scenario is shown in grey and the no-viaduct scenario in red.




Figure 9: Posterior predictive distributions of average travel times in 2020 for additional routes that directly involve the viaduct. The posterior distributions are represented by boxplots, with the box containing the interquartile range with the median marked, and the whiskers covering the 95% posterior confidence interval. The baseline scenario is shown in grey and the no-viaduct scenario in red.




Figure 10: Posterior predictive distributions of differences between average travel times for the two scenarios for all routes. The difference are equal to average travel time for the no-viaduct scenario minus that for the baseline (viaduct) scenario. The posterior distributions are represented by boxplots, with the box containing the interquartile range with the median marked, and the whiskers covering the 95% posterior confidence interval. The routes above the horizontal line are those shown in Figure 4 that do not include the viaduct, while the routes below the line are those that do contain the viaduct.

for route 13 falls outside our 95% confidence interval.

We should make clear that our analysis is not directly comparable with the WSDOT study. Not only are the prediction years different, but more importantly, the WSDOT study does not consider the long-term changes that occur when transportation projects change patterns of accessibility: households can relocate to be closer to their jobs, workers can change jobs to reduce their commute, businesses can relocate to take advantage of better access at different locations, and real estate developers can respond to new opportunities to develop housing and non-residential space. Our analysis integrates these forms of long-term adaptation, in addition to the short-term adaptations that travelers have when accessibility patterns change: they can change destinations, times of travel, modes of travel, and routes. The collection of these kinds of adaptive behaviors provides a reservoir of flexibility that has not previously been thoroughly examined. In instances such as a temporary or even long-term closure of a major transportation facility, the reality in terms of traffic conditions is often far better than transportation officials expect. These kinds of adaptive behaviors provide a plausible explanation, though many other factors could also contribute.

What our results suggest, in short, is that even using a worst-case scenario and comparing it to a capacity-neutral replacement of the Alaskan Way Viaduct, the travel time benefits of the higher capacity alternative are modest, and fairly localized to the viaduct corridor. There does not appear to be much effect on longer commutes or on I-5 in the vicinity of downtown, as evidenced by the overlapping distributions of the predicted travel times. Further, our combined analysis of land use and transportation reveals considerably more adaptive capacity than the analysis done by the WSDOT, which considers only travel changes and excludes by assumption any adaptation in location choices of households, firms and real estate development. Accounting for uncertainty, in short, the expectations of benefits from maintaining the current level of traffic capacity in the viaduct corridor may be higher than can be scientifically supported by the available models and evidence.

In future research, we hope to further develop the methodology described in this paper, and to incorporate refinements in the current generation of models used in supporting the decision-making process on large-scale infrastructure projects such as the Alaskan Way Viaduct.

7 Acknowledgments

This work has been funded by NSF grants EIA-0121326 and IIS-0705898, and EPA grant R831837. The authors would like to thank Larry Blain for helpful discussions and for helping to configure the travel model, and the Puget Sound Regional Council for providing access to the travel model and data.

References

- Boyce, D. (1984). Urban transportation network-equilibrium and design models: Recent achievements and future prospects. *Environment and Planning A* 16, 1445–1474.
- Clay, M. and R. Johnston (2006). Multivariate uncertainty analysis of an integrated land use and transportation model: Meplan. *Transportation Research Part D* 11(3), 191 – 203.
- Downs, A. (2004). *Still Stuck in Traffic: Coping with Peak-hour Congestion*. Washington, D.C.: The Brookings Institution.
- Gilmore, S. (2008). New viaduct speediest option. *Seattle Times*, Local News, 11/14/2008. Available at http://seattletimes.nwsourc.com/html/localnews/2008389297_viaduct14m.html.
- Kockleman, K. (2002). Uncertainty propagation in an integrated land use-transport modeling framework: Output variation via UrbanSim. *Transportation Research Record* (1805), 128–135.
- Kockleman, K. (2003). Propagation of uncertainty in transportation-land use models: An investigation of DRAM-EMPAL and UTPP predictions in Austin, Texas. *Transportation Research Record* (1831), 219–229.
- McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers in Econometrics*, pp. 105–142. New York: Academic Press.
- McFadden, D. (1978). Modeling the choice of residential location. In A. Karlqvist, L. Lundqvist, F. Snickars, and J. Wiebull (Eds.), *Spatial Interaction Theory and Planning Models*, pp. 75–96. Amsterdam: North Holland.

- McFadden, D. (1981). Econometric models of probabilistic choice. In C. Manski and D. McFadden (Eds.), *Structural Analysis of Discrete Data with Econometric Applications*, pp. 198–272. Cambridge, MA: MIT Press.
- Poole, D. and A. E. Raftery (2000). Inference for deterministic simulation models: the Bayesian melding approach. *Journal of the American Statistical Association* 95(452), 1244–1255.
- Raftery, A. E., G. H. Givens, and J. E. Zeh (1992). Inference from a deterministic population dynamics model about bowhead whale, *balaena mysticetus*, replacement yield. Technical Report 232, Department of Statistics, University of Washington, Seattle, Wash.
- Raftery, A. E., G. H. Givens, and J. E. Zeh (1995). Inference from a deterministic population dynamics model for bowhead whales. *Journal of the American Statistical Association* 90(430), 402–416.
- Ševčíková, H., A. Raftery, and P. Waddell (2007). Assessing uncertainty in urban simulations using Bayesian melding. *Transportation Research B* 41, 652–669.
- Sumalee, A., P. Luathep, W. H. K. Lam, and R. D. Connors (2009). Transport network capacity evaluation and design under demand uncertainty. *Transportation Research Record*.
- Waddell, P. (2002). UrbanSim: Modeling urban development for land use, transportation, and environmental planning. *Journal of the American Planning Association* 68(3), 297–314.
- Waddell, P., A. Borning, M. Noth, N. Freier, M. Becke, and G. Ulfarsson (2003). Microsimulation of urban development and location choices: Design and implementation of UrbanSim. *Networks and Spatial Economics* 3(1), 43–67.
- Waddell, P., G. Ulfarsson, J. Franklin, and J. Lobb (2007). Incorporating land use in metropolitan transportation planning. *Transportation Research Part A: Policy and Practice* 41, 382–410.
- WSDOT (2004). Washington state department of transportation: Sr 99 - alaskan way viaduct and seawall replacement. Web page: <http://www.wsdot.wa.gov/projects/Viaduct>.
- WSDOT (2005). Washington state department of transportation: About the alaskan way viaduct and seawall. Web page: <http://www.wsdot.wa.gov/Projects/Viaduct/viaductSeawall.htm>.
- WSDOT (2007). Washington state department of transportation: Alaskan way viaduct and seawall replacement - library - meeting materials. Web page: <http://www.wsdot.wa.gov/Projects/Viaduct/library-meetingmaterials.htm>.
- WSDOT (2008). Washington state department of transportation: Transportation results present different choices after alaskan way viaduct begins to come down in 2012. Web page: <http://www.wsdot.wa.gov/News/2008/11/13-AWVchoices.htm>.