

Assessing Uncertainty in Simulation Based

Maritime Risk Assessment

Abstract

Recent work in the assessment of risk in maritime transportation systems has used simulation-based probabilistic risk assessment techniques. In the Prince William Sound and Washington State Ferries risk assessments, the studies' recommendations were backed up by estimates of their impact made using such techniques and all recommendations were implemented. However, the level of uncertainty about these estimates was not available, leaving the decision-makers unsure whether the evidence was sufficient to assess specific risks and benefits. The first step towards assessing the impact of uncertainty in maritime risk assessments is to model the uncertainty in the simulation models used. In this paper, a study of the impact of proposed ferry service expansions in San Francisco Bay is used as a case study to demonstrate the use of Bayesian simulation techniques to propagate uncertainty throughout the analysis. The conclusions drawn in the original study are shown, in this case, to be robust to the inherent uncertainties. The main intellectual merit of this work is the development of Bayesian simulation technique to model uncertainty in the assessment of maritime risk. However, Bayesian simulations have only been implemented as theoretical demonstrations. Their use in a large, complex system may be considered state of the art in the field of computational sciences.

Keywords: *Uncertainty Analysis, Bayesian Simulation, Maritime Risk Assessment.*

1. Introduction

The grounding of the *Exxon Valdez* in Prince William Sound, the capsizing of the *Herald of Free Enterprise* and the *Estonia* passenger ferries are some of the most widely publicized accidents in marine transportation systems. The consequences of these accidents range from severe environmental damage to large-scale loss of life, which leads to the immediate questions of how to prevent such accidents in the future and how to mitigate their consequences if they should occur.

The presence of uncertainty in analyzing risk is well recognized and discussed in the literature. However, these uncertainties are often ignored or under-reported in studies of controversial or politically sensitive issues [1]. Two types of uncertainty discussed in the literature are aleatory uncertainty (the randomness of the system itself) and epistemic uncertainty (the lack of knowledge about the system). In a modeling sense, aleatory uncertainty is represented by probability models that give probabilistic risk analysis its name [2], while epistemic uncertainty is represented by lack of knowledge concerning the parameters of the model [3]. In the same manner that addressing aleatory uncertainty is critical through probabilistic risk analysis, addressing epistemic uncertainty is critical to allow meaningful decision-making. Several examples have been published of the conclusions of an analysis changing when uncertainty is correctly modeled [4]. There is one other type of uncertainty and that is uncertainty about the model itself [5].

Previous work in maritime risk assessment has considered aleatory uncertainty, but not epistemic uncertainty nor model uncertainty. Early work in maritime risk assessment concentrated on assessing the safety of individual vessels or marine structures, such as nuclear powered vessels [6], vessels transporting liquefied natural gas [7] and offshore oil

and gas platforms [8]. Later, the USCG attempted to prioritize federal spending to improve port infrastructures using a classical statistical analysis of nationwide accident data [9,10]. More recently, quantitative risk assessment (QRA) has been introduced in the assessment of risk in the maritime domain [11-18]. Pate-Cornell [1] defines six levels of treatment of uncertainty in risk analysis: 0. Identification of hazards; 1. Worst case analysis; 2. Plausible upper bound analysis; 3. Best estimates; 4. Probability and risk analysis; 5. Display of risk uncertainties. On this scale, the early work in maritime risk assessment could be classified as level 3, providing best estimates of accident risk, while the later work using QRA can be classified as level 4. None of the work has considered epistemic uncertainty or the dynamic nature of the aleatory uncertainty.

In a maritime transportation system (MTS) traffic patterns change over time in a complex manner causing inherent aleatory uncertainty. System simulation has been proposed to model aleatory uncertainty in analyzing port operations [19-21]. In addition to the dynamic nature of traffic patterns, situational variables such as wind, visibility and ice conditions change over time making risk a dynamic property of the system. The Prince William Sound (PWS) Risk Assessment [22,23] and the Washington State Ferries (WSF) Risk Assessment [24] differ from previous maritime risk assessments as the dynamic nature of risk was captured by integrating system simulation [25] with available techniques in the field of QRA [2] and expert judgment elicitation [26]. The PWS and WSF Risk Assessments were conducted at level 4 on the scale of Pate-Cornell [1], considering aleatory uncertainty in a more explicit manner than non-simulation based approaches. However, epistemic uncertainty was not considered.

The simulation in a maritime risk assessment model is used to estimate the frequency of the various possible system states. Random processes control the arrival of vessels and environmental conditions in to the simulation. The parameters of the random processes must be estimated from data collected from the system. This introduces inherent epistemic uncertainty in the inputs to the simulation. There is also model uncertainty in selecting appropriate random arrival processes. Output estimates of from simulations are also uncertain as the simulation cannot be run infinitely many times or for an infinite run time. Thus there is output uncertainty, both epistemic and aleatory.

While epistemic uncertainty can be addressed through frequentist statistical techniques such as bootstrap or likelihood based methods [27], the Bayesian paradigm is widely accepted as a method for dealing with both types of uncertainty [26,28,29,30]. Bayesian modeling can allow for the distinction and handle the underlying differences inherently when used for analyzing data and expert judgments [31-35]. To move to level 5, epistemic uncertainties must be addressed in maritime risk assessment and we must use Bayesian analytical methods in building and analyzing our simulation models. There is an extensive literature on the theory of Bayesian simulation analysis [36-42]. The general paradigm of this research is that simulation analysis is a decision-making tool and therefore should be used under a decision-analytic framework. Barton and Schruben [43] show several examples of the conclusions of simulation studies changing when input uncertainty is incorporated rather than using mean input estimates. Chick [44] gives an excellent review of Bayesian methods, discussing both input and output uncertainty. Chick also discusses model uncertainty. However, due to the computational complexity of a Bayesian simulation of a port system with many different types of vessels on many

different routes, the scope of the development herein is restricted to aleatory and epistemic uncertainty. Handling model uncertainty was simply not computationally feasible at this point.

The project discussed herein is funded by National Science Foundation (NSF) grants SES 0213627 and SES 0213700. To complete the objective of a complete theoretical framework for modeling uncertainty in full-scale risk assessments, such as that performed for PWS or the WSF, the following tasks need to be completed.

- Task 1: Representation of Uncertainty in the Simulation
- Task 2: Representation of Uncertainty in the Expert judgment
- Task 3: Propagation of the Uncertainties Through the Whole Model
- Task 4: Conduct a Trial Uncertainty Analysis

This paper discusses the work on Task 1, specifically the incorporations of uncertainties in the simulation part of the overall model through Bayesian simulation techniques.

The content of this paper is as follows. Section 2 outlines a study of the San Francisco Bay ferries. We will use this study to demonstrate the developments proposed herein. The modeling of input uncertainty is discussed in Section 3. A method for modeling output uncertainty is discussed in Section 4 and the format of the display of these results is outlined in Section 5. These techniques are used to create a Bayesian simulation of the San Francisco Bay maritime transportation system in Section 6. The impact of epistemic uncertainty on the results for the current ferry system and the comparison of the expansion alternatives are discussed in Sections 7, 8 and 9. Conclusions and areas for future research are given in Section 10.

2. Case Study: Expansion of the San Francisco Bay Ferries

As an example of the application of uncertainty analysis in a maritime simulation model, we will extend an analysis performed for the San Francisco Bay Water Transit Authority [45]. In an effort to relieve congestion on freeways, the state of California is proposing to expand ferry operations on San Francisco Bay. The three proposed expansion scenarios are: Alternative 3: Enhanced Existing System; Alternative 2: Robust Water Transit System and Alternative 1: Aggressive Water Transit System. From these, Alternative 3 is the least aggressive expansion scenario and Alternative 1 is the most aggressive one. As part of the analysis, a simulation of the maritime transportation system in San Francisco Bay was used to assess the effect on the level of vessel interactions of proposed changes [44]. A simulation model was created capable of estimating the increase in the number of vessel interactions in the current system as well as in three alternative expansion plans. Figure 1 shows a snapshot of San Francisco Bay in the simulation. For a more detailed look, movies of the simulation for each of the cases can be viewed at <http://www.people.vcu.edu/~jrmerric/SFBayMovies/>.

Due to time and budget constraints a full-scale risk assessment, such as the previous work in the PWS or WSF Risk Assessments, was not feasible. Instead, to assess the impact of aggressive ferry expansion, the scope of the San Francisco Bay study was limited to the simulation part of the model, leaving the accident probability part to a later project if the expansion proposal is approved. The output of the model is a map showing the frequency of interactions across the study area, representing the level of congestion under each alternative. Figure 2 shows the results of the simulation for the current ferry service.

This analysis models the aleatory uncertainty about the traffic patterns through the simulation model, but does not address sources of epistemic uncertainty, making it a level 4 analysis on Pate-Cornell's scale. However, the analysis of the current ferry service and a comparison of four proposed alternatives were submitted to the legislature as part of the overall analysis and will be used in the expansion decision. The incorporation of uncertainty in the simulation of San Francisco Bay will allow the implementation of the framework developed in a situation where the results will have meaning without having to first complete the other tasks in the project. While the main intellectual merit of this work is the development of an overarching framework for including uncertainty in the assessment of maritime risk, Bayesian simulation analysis techniques have only been proposed in theoretical settings, thus their use in a large complex system may be considered state of the art in the field of computational sciences.

3. Modeling Input Uncertainty

Input uncertainty should be incorporated in the analysis to reflect the limited data available to populate the parameters of the arrival processes in a simulation model [41]. In the San Francisco Bay study, if we consider traffic arrivals to the system only, there were 5,277 separate arrival processes for various types of vessels and routes [45]. These arrival processes can be modeled by the standard renewal process [46], with a probability distribution chosen to model the inter-arrival times. Historical inter-arrival times are calculated from data supplied by the Vessel Traffic Service on Treasure Island. Let $T_1^k, \dots, T_{m^k}^k$ be the m^k independent inter-arrival times for the k -th arrival process ($k = 1, \dots, 5277$).

In a classical simulation approach, the probability model is usually chosen by determining best estimates of the parameters from the data for several possible families of distributions and comparing the fit of each distribution to the data using fit statistics such as the Andersen Darling, Chi-square or Kolmogorov-Smirnov statistics [46]. Suppose $F_1^k(t | \Theta_1^k), \dots, F_p^k(t | \Theta_p^k)$ are p families of probability distribution, such as the exponential, Weibull, gamma or log-normal distributions. The superscript k is included throughout as each arrival process can be modeled by a different probability distribution and will certainly have different parameter values. Best estimates of each set of parameters, $\hat{\Theta}_j^k$, are obtained from the data $D^k = \{T_1^k = t_1^k, \dots, T_{m^k}^k = t_{m^k}^k\}$, using maximum likelihood, method of moments or other estimation procedures. The best fit distribution is then chosen by taking either the fitted distribution with the lowest appropriate fit statistic or at least a fitted distribution that is not rejected by the corresponding hypothesis test and that has desirable properties, such as simple manipulation of the mean or variance. Thus aleatory uncertainty is modeled by the renewal process, but as only best estimates are used for the parameters of the generating probability distribution, the uncertainty about their true values is not included in the model.

Under the Bayesian paradigm, prior distributions are specified for the parameters of the postulated distributions, denoted by $\pi_1^k(\Theta_1^k), \dots, \pi_p^k(\Theta_p^k)$, and the data is used to update these priors using the standard Bayesian machinery to obtain posterior distributions denoted by $\pi_1^k(\Theta_1^k | D^k), \dots, \pi_p^k(\Theta_p^k | D^k)$. To demonstrate, Bayesian updating procedures let us consider container ships arriving from an offshore anchor point passing under the Golden Gate Bridge and birthing in the Oakland Outer Harbor. Overall 176 such transits occurred from 7/31/1998 and 12/31/2001, with an average of

4.44 days between transits. We will consider the exponential distribution. For the exponential distribution with parameter λ^k , the gamma distribution is a natural conjugate prior for λ^k . That is, if λ^k is assumed a priori to be drawn from a gamma distribution with shape parameter a^k and scale parameter b^k , then after updating with the inter-arrival time data, λ^k will be gamma distribution with shape parameter $a^k + \sum_1^{m^k} t_i^k$ and scale parameter $b^k + m^k$. For our container route, we assume a vague prior by setting $a^k = 0.001$ and $b^k = 0.001$, which corresponds to a prior mean of 1 and a prior variance of 1000. For this route, $\sum_1^{m^k} t_i^k = 781.44$ and, as previously mentioned, $m^k = 176$, thus a posterior, $a^k = 781.441$ and $b^k = 176.001$. In the simulation, inter-arrival times for this process could then be sampled by first sampling from a gamma distribution with shape 781.441 and scale 176.001 to obtain a sample for λ^k and then sampling from an exponential distribution with the parameter set to the sampled value of λ^k (Chick 2000). Equivalently, inter-arrival times could be sampled from a Pareto distribution with shape 781.441 and scale 176.001 [2, ch. 4].

The most difficult problem in the Bayesian approach is choosing the best fitting probability model. Recent work in the field of Bayesian statistics has included criteria such as Bayes factors [47], posterior predictive densities [48] and the recently proposed Decision Information Criterion [49]. We will demonstrate the use of the Decision Information Criterion (DIC). The Bayesian deviance is defined as

$$D(\Theta_j^k) = -2 \ln p(D^k | \Theta_j^k) + 2 \ln f(D^k)$$

where $f(D^k)$ is some fully specified standardizing term that is a function of the data alone and thus does not affect the model comparison. The model fit is then represented by $\bar{D} = E[D(\Theta_j^k) | D^k]$, the expected Bayesian deviance after updating with the available data. An estimate of the effective number of parameters is given by $p_D = E[D(\Theta_j^k) | D^k] - D(E[\Theta_j^k])$, the difference between the expected Bayesian deviance after updating with the available data and the Bayesian deviance calculated at the expected value of the parameters after updating with the available data. The DIC is then equal to $\bar{D} - p_D$, the model fit penalized by the number of parameters of the model.

To demonstrate this approach to the choice of probability distribution, we will compare the exponential, Weibull, gamma and log-normal distributions for the arrival process discussed above, with appropriate vague priors chosen for the parameters of each distribution. Table I shows the DIC results. The calculations in Table 1 were performed in WinBugs version 1.4 [50]. Notice that the effective number of parameters is quite close to the true number of parameters in the model, one for the exponential distribution and two for the rest. Overall, the gamma distribution has the best ranking, although the difference with the Weibull is negligible and could be explained by sampling error.

The reader should note that we have not fully considered model uncertainty in our approach. The “winner” in the DIC competition is the only model used in the simulation even though the DIC statistic is based on finite data and so the “true” model is not known. Other approaches would be to use a mixture of potential inter-arrival time distributions [5] or to use a non-parametric choice, such as a Dirichlet process on the cumulative distribution function [51]. While more complete in their treatment of model

uncertainty, such approaches are simply not computationally feasible with such a large-scale simulation model. Hence our less computational approach.

The next question is how to propagate this input uncertainty through the simulation to the outputs of interest. The sampling algorithm can be stated as follows [40,44]:

For $r=1, \dots, n$ replications:

1. Sample values of Θ_*^k , the parameters of the chosen probability model conditioned on the available arrival data, for each arrival process ($k = 1, \dots, 5277$) to be used in the r -th replication,.
2. For the r -th replication:
 - a. Sample random variates for the inter-arrival times given the parameter values drawn for each Θ_*^k ;
 - b. Generate the simulation output that is determined by these random variates.

End loop

Thus for our exponential distribution example, we sample values for each λ^k from the corresponding gamma distributed posterior distribution given the inter-arrival time data. These values are then used throughout a replication to sample from the exponential inter-arrival time probability model. New values for each λ^k are used for each replication. We should note that this implies that each value of the outputs of interest represents only aleatory uncertainty; to examine epistemic uncertainty, one must run a sufficient number of replications to cover the range of the posterior distributions on the parameters of the input probability models.

4. Modeling Output Uncertainty

In our risk assessment methodology, the quantity of interest is the yearly number of vessel interactions; the data obtained from the simulation in each replication will be the number of vessel interactions occurring in each replication of the simulation, denoted N_r , for the r -th replication. In the analysis of output data, the focus of Bayesian simulation research has been on estimating means of important output statistics, rather than attempting to define its probability distribution. Bayesian Model Averaging is the commonly used term when the average of the s output statistics obtained from the simulation is used to estimate the statistic's expected value.

However, as we wish to propagate uncertainty throughout the model, a probability model will be hypothesized for the output statistic. The output values for the replications of the simulation are treated as data to update the prior distributions on the output statistic model's parameters. Chick [41] notes that this can be thought of as a Bayesian version of metamodeling [46]. Such treatment of output data flows naturally in to a decision-analytic handling of choosing the best system [40, 42, 44].

As our output data is in the form of a count, the number of vessel interactions can be naturally modeled using a Poisson distribution with rate μ , with a conjugate gamma distributed prior on μ with shape α and scale β . The likelihood function for

s replications of the simulation, $L(\mu | N_1 = n_1, \dots, N_s = n_s) = \prod_{i=1}^s \frac{\mu^{n_i}}{n_i!} e^{-\mu}$, is used to update the prior (usually a vague prior) with the simulated data. The posterior distribution of the vessel interaction frequencies will be a gamma distribution with shape $\alpha + \sum_{i=1}^s n_i$ and

scale $\beta + s$. This distribution includes the epistemic uncertainty about the expected number of interactions, while the Poisson distributed probability model represents the aleatory uncertainty about the actual number of interactions in a given year.

5. Format of the Results

There are several useful ways of portraying the information contained in this analysis. We will depict the frequency of vessel interactions each alternative in aggregate as well spatially across the study area for. However, we will also depict the differences between the alternative ferry systems in the alternatives. To accomplish this aim, one can compare the distribution of the expected number of interactions in a given period across different geographic locations in the study area or across different alternative simulations. These distributions will incorporate the aleatory uncertainty as they are obtained from the simulation and the epistemic uncertainty as Bayesian simulation was used.

We create maps of the quantiles of the expected rate of interactions in each of a grid of cells across the San Francisco Bay. These maps were the main output format used in the original study as they allow decision makers to assess the risk across the system (Figure 2); the addition of uncertainty through quantile maps allows decision-makers to assess the impact of uncertainties on the conclusions they draw from these maps. If we denote $\mu^a[x, y]$ as the expected number of interactions in the grid cell indexed by x and y for alternative a , then we will create maps of the 5th, 50th and 95th percentiles of the posterior distribution of the $\mu^a[x, y]$'s for each alternative.

Other comparisons are possible comparing the current ferry system, or Base Case, to the three proposed expansion alternatives. This comparison can be made in aggregate

using the posterior distribution of the yearly total expected number of interactions in the whole system for each alternative. However, more detail can be obtained by comparing the yearly expected number of interactions in each of the grid cells across the San Francisco Bay. We will calculate the probability that the rate in a given grid cell in one alternative is greater than or equal to that for the same cell in another alternative. Following the above notation, we wish to calculate $P(\mu^a[x, y] > \mu^b[x, y])$ for all grid cells indexed by x and y and all combinations of alternatives a and b . These maps can be called probability dominance maps and can give decision-makers a great deal of information as they not only indicate when one alternative is likely to have more interactions, but also the level of certainty in this assertion. As $P(\mu^a[x, y] > \mu^b[x, y])$ cannot be obtained in closed form for the gamma distribution, we use sampling approximations by sampling iteratively from $\mu^a[x, y]$ and $\mu^b[x, y]$ and calculating the proportion where $\mu^a[x, y] > \mu^b[x, y]$.

6. A Bayesian Simulation of San Francisco Bay

We created a simulation of the San Francisco Bay maritime transportation system using the Bayesian approach to input analysis by modifying the program used in the original study [45]. All non-ferry traffic, except scheduled regattas, was modeled in the manner discussed in the previous section. As there are 5,277 arrival processes and as this is a demonstration, we chose the exponential distribution to model the inter-arrival times for each process. As the sufficient statistics necessary to perform such an update are $\sum_1^{m^k} t_i^k$ and m^k , we could perform a database query on the San Francisco Bay Vessel Traffic

Service's log of transits to obtain these quantities. As the sum of the inter-arrival times is the same as the time between the first and last arrival in the database, the query returned the first arrival, the last arrival and the total number of log entries for each combination of vessel type, origin and destination. Thus the posterior distribution of the rate of arrivals could be easily obtained.

To implement our approach within the simulation, code was added to the simulation that sampled from the posterior distribution of each λ^k prior to each replication of a year. Within each replication, samples were taken for the inter-arrival times from exponential distributions with parameters set to the sampled values for λ^k . In the existing simulation, the ferry transits for the current ferry system and each of the alternatives were based on a fixed schedule. Visibility and wind conditions were incorporated by tracing large databases of environmental data obtained from National Oceanographic and Atmospheric Administration (NOAA) observation stations in the study area. The vessel interaction counting methodology was also programmed in to the simulation. For further details of these existing pieces of the model, we refer the reader to our previous work [45].

Within the simulation, a snapshot of the system is calculated for each simulated minute based on the previous snapshot, randomly drawn new vessel arrivals and the environmental conditions models. All ferries in the system are considered one at a time to search for possible interactions with collision potential. For each ferry, the track of the ferry is compared to the track of each other vessel in the simulation and the time to cross is calculated. If the crossing occurs close enough to the ferry and will occur in the next 15

minutes, an interaction between the two vessels is counted. These counts are then used to update the output meta-model.

One drawback of this simulation-based approach is its heavy computational load. Running 50 years of Alternative 3 took 24 hours on a high-end, multi-processor workstation. Even 10 years of Alternative 1 took 30 hours, so a full 50 years would have taken almost 6 days. While snapshot calculation time increases linearly with the number of vessels in the simulation, the search for interactions increases exponentially with the number of ferries as each ferry's track must be compared to each other vessel's track, including other ferries. In the later stages of a risk assessment, multiple such runs must be run for a complete analysis. With such long run times, this is not currently feasible. We intend to study the possibilities for spreading the computational load across a network of less expensive computers. This will allow the implementation of these techniques without prohibitive hardware costs. However, the current simulation code is simply made as efficient as possible with careful programming. For instance, the program uses databases for the ferry schedules performing a query for each simulated day's schedule of transits. This makes the schedule determination more efficient. Results are accumulated in internal arrays and are only written to output files at the end of the run. However, there are no special tricks and no tasks that can be exploited with efficient algorithms.

7. Uncertainty Results for the Current Ferry System

The simulation was run for one replication of one day of the current ferry system, or Base Case, specifically July 3rd, a moderately busy Monday in the summer. We assumed a vague prior for the expected number of interactions in each grid cell, $\mu^a[x, y]$, by setting each $\alpha = 0.001$ and $\beta = 0.001$, which corresponds to a prior mean of 1 and a prior

variance of 1000 . In total there were 9,430 interactions in the simulated day. The posterior distribution of each $\mu[x, y]$ was calculated and summed over all x and y to find the posterior distribution of the total expected number of interactions. This distribution has a median of 9,430 interactions, a 5th percentile of 6802 interactions and a 95th percentile of 12,711 interactions.

To reflect the results across the grid of cells, 5th, 50th and 95th percentile maps were created (Figure 3). The box at the center of the map surrounds the ferry building in San Francisco. The accompanying count shows the percentage of the interactions that occur in this vicinity. One can see that with each replication simulating one day, there is considerable variability about this quantity. However, examining the color of the cells, with darker cells having more interactions, the colors do not change much from the 5th to the 95th percentile. There is a small change in the interactions on the northerly route, but the colors in the central bay area do not change. Thus any conclusions drawn from these maps are robust to the uncertainties in the simulation even with only one day simulations.

Obviously, in such a risk assessment we do not perform such small numbers of replications. Furthermore risk can change throughout the year due to environmental or traffic pattern changes. Thus a full year of simulation is considered one replication and the quantity of interest is the expected yearly number of interactions. We performed 100 replications of a year finding an average of 640,889 interactions per year. The posterior distribution of the total yearly expected number of interactions has a median of 640,889 interactions, a 5th percentile of 640,417 interactions and a 95th percentile of 641,359 interactions, indicating a small range of uncertainty. Figure 4 shows the 5th, 50th and 95th percentile maps for the 100 replications of a year. There is very little variability either in

the count of interactions in the ferry building area or in the colors of the grid cells. Thus all conclusions drawn from the Base Case analysis are robust to the inherent uncertainty.

8. Uncertainty Results for the Aggregate Alternatives

Comparison

100 replications of one day were simulated for the base case and each of the three alternatives to obtain the posterior distribution of the number of interactions in the simulated day (again July 3rd). Figure 5 shows the comparison, plotting the median of this distribution against the total number of ferry transits in each simulation. Error bars are also added to indicate the range from the 5th to the 95th percentiles of the posterior distribution for each alternative. A major conclusion drawn from the original study in San Francisco Bay was that the number of ferry to vessels interactions grows exponentially with the number of ferry transits, not linearly, and thus the safety levels currently enjoyed by the San Francisco Bay ferry service cannot be maintained under the planned expansion scenarios without equally aggressive investment in risk intervention [44]. Figure 5 shows that this conclusion is not affected by the epistemic uncertainties in the results. Despite the replication of just one day for each alternative, the level of uncertainty in these posterior distributions is small relative to the large differences between the alternatives.

Figure 6 reinforces this conclusion with 100 replications of a year for the Base Case and Alternative 3 and 25 replications of a year for the Alternatives 2 and 1 (due to the run time). The same error bars are included but are not visible on this scale. In the original study, we concluded that a linear increase in the number of ferry transits in the bay will

lead to an exponential growth in the number of interactions. This result shows that the conclusion is robust to the inclusion of epistemic uncertainty in the modeling process.

9. Uncertainty Results for the Geographic Alternatives

Comparison

The aggregate analysis shows conclusive differences between the current ferry system and the three alternatives. Each addition of ferry transits results in additional interactions, with the growth being exponential not linear. However, is there more to these comparisons. Are these additions in certain areas or hot spots? Do some areas actually see fewer interactions? Examination of the probability dominance maps can give a more detailed picture of the differences between the alternatives.

Figure 7 shows the probability dominance map for the Base Case compared to Alternative 3. These maps are easier to read in color, but are converted to grayscale here for reproduction. As indicated in the legend, black cells indicate almost certainty that Alternative 3 will see more interactions in that location than the Base Case, with less certainty shown in dark gray. Almost white cells show the reverse with almost certainty, while light gray indicates less certainty. The numbers of interactions in medium gray cells are not different between these two scenarios. That is, the posterior distributions of the expected number of interactions in a medium gray cell are almost identical between the two alternatives mapped. Thus, the probability that one is higher than the other is 0.5 (50%) and corresponds to medium gray on the color scale.

The majority of the grid cells in Figure 7 are black, reinforcing the conclusions from the aggregate results that Alternative 3 has significantly more interactions overall than the Base Case. However, there are some almost white cells showing the reverse

conclusion. The main area of white is around the Golden Gate Bridge and Richardson Bay. The ferries in this area are running from San Francisco to Sausalito and Tiburon or are tours around the Bay visiting the Golden Gate Bridge. The tours were unchanged from the Base Case to the alternatives. However, one problem with the schedules supplied for the alternatives was that they consisted of a start time, end time and time between ferries. For the Sausalito and Tiburon ferries, they start at 7 am and run every 30 minutes until 10 pm during the week. At the weekend they run every 60 minutes. This is significantly more than in the Base Case, but this means that there are definite patterns to the transits that are not reflective of a more mature schedule. These ferries do not interact as much because of the timing of the transits. However, the number of interactions in this area is very low relative to other areas of the Bay. This is also shown on the northerly routes to Larkspur and Vallejo. The almost white in the middle of the black area is actually on the ferry routes. There are only certain places where ferries going in different directions meet due to the schedule (the black cells along the route). In other parts of the route fewer interactions occur (almost white cells).

Figure 8 shows another comparison of interest, Alternative 2 versus Alternative 1. This map is easier to understand. Alternative 1 has a number of additional routes that run the length of the study area from the northeast to the south. Along the center of the navigable area there are more interactions in Alternative 1, indicated by the large areas of black. Around the black cells there are some thin bands of dark gray, indicating some uncertainty. However, away from the center and along routes that remain the same between the two alternatives, the numbers of interactions are the same, indicated by

medium gray cells. Around the edges other colors appear, but this is a result of a very small number of interactions even over the 10 replications of a year.

10. Conclusions

A Bayesian approach to simulation modeling has allowed the treatment of epistemic uncertainty concerning the movements of non-ferry traffic as well as the aleatory uncertainty captured by the simulation model itself. We reviewed a Bayesian approach to analyzing input data and developed a Bayesian meta-model that allowed for interesting and useful output formats depicting the output uncertainty. In particular, we used this approach to examine the impact of uncertainty on the conclusions drawn in a study of proposed ferry service expansions in San Francisco Bay. We used maps of the percentiles of the posterior distribution of the expected number of interactions across the study area to show that conclusions drawn from these geographic profiles of vessel interactions are robust to the inherent uncertainties. Further maps were developed that showed the probability that one alternative would have more interactions than another over the study area. These maps allowed for detailed comparisons of the alternatives.

As mentioned in the introduction, the work discussed herein represents the first task necessary for a full-scale maritime risk assessment considering both aleatory and epistemic uncertainty. The next task is to develop a Bayesian accident probability model incorporating historical accident and incident data and expert judgments. The accident probability model must then be integrated with the output meta-model from the simulation for full scale risk results, another highly computational task. This work is forthcoming.

Disclaimer

This material is based upon work supported by the National Science Foundation under Grant Nos. SES 0213627 and SES 0213700. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. We thank the Area Editor and two anonymous reviewers for their helpful comments that have improved both the analysis and the description of this work.

References

- [1] Paté-Cornell, M. E. 1996. Uncertainties in risk analysis: six levels of treatment. *Reliability Engineering and System Safety* **54**(2-3) 95-111.
- [2] Bedford, T. M., R. M. Cooke. 2001. *Probabilistic Risk Analysis: Foundations and Method*. Cambridge University Press, Cambridge UK.
- [3] Parry, G. 1996. The characterization of uncertainty in Probabilistic Risk Assessments of complex systems. *Reliability Engineering and System Safety* **54** 119-126.
- [4] Cooke, R. 1997. Uncertainty modeling: examples and issues. *Safety Science* **26**(1) 49-60.
- [5] Zio, E., G. E. Apostolakis, 1996. Two methods for the structured assessment of model uncertainty by experts in performance assessments of radioactive waste repositories. *Reliability Engineering and System Safety* **54** 225-241.
- [6] Pravda, M. F., R. G. Lightner. 1966. Conceptual study of a supercritical reactor plant for merchant ships. *Marine Technology* **4** 230-238.

- [7] Stiehl, G. L. 1977. Prospects for shipping liquefied natural gas. *Marine Technology* **14**(4) 351-378.
- [8] Paté-Cornell, M. E. 1990. Organizational aspects of engineering system safety: The case of offshore platforms. *Science* **250**(4985) 1210-1217.
- [9] U.S. Coast Guard. 1973. *Vessel Traffic Systems: Analysis of Port Needs*. Report No. AD-770 710. Washington, D.C.: U.S. Coast Guard.
- [10] Maio, D., R. Ricci, M. Rossetti, J. Schwenk, T. Liu. 1991. *Port Needs Study*. Report No. DOT-CG-N-01-91-1.2. Prepared by John A. Volpe, National Transportation Systems Center. Washington, D.C.: U.S. Coast Guard.
- [11] Hara, K., S. Nakamura. 1995. A comprehensive assessment system for the maritime traffic environment. *Safety Science* **19**(2-3) 203-215.
- [12] Roeleven, D., M. Kok, H. L. Stipdonk, W. A. de Vries. 1995. Inland waterway transport: modeling the probabilities of accidents. *Safety Science* **19**(2-3) 191-202.
- [13] Kite-Powell, H. L., D. Jin, N. M. Patrikalis, J. Jebsen, V. Papakonstantinou. 1996. *Formulation of a Model for Ship Transit Risk*. MIT Sea Grant Technical Report, Cambridge, MA, 96-19.
- [14] Slob, W. 1998. Determination of risks on inland waterways. *Journal of Hazardous Materials* **61**(1-3) 363-370.
- [15] Fowler, T. G., E. Sorgard. 2000. Modeling ship transportation risk. *Risk Analysis* **20**(2) 225-244.
- [16] Trbojevic, V. M., B. J. Carr. 2000. Risk based methodology for safety improvements in ports. *Journal of Hazardous Materials* **71**(1-3) 467-480.

- [17] Wang, J. 2000. A subjective modeling tool applied to formal ship safety assessment. *Ocean Engineering* **27**(10) 1019-1035.
- [18] Guedes Soares, C., A. P. Teixeira. 2001. Risk assessment in maritime transportation. *Reliability Engineering and System Safety* **74**(3) 299-309.
- [19] Andrews, S., F. H. Murphy, X. P. Wang, S. Welch. 1996. Modeling crude oil lightering in Delaware Bay. *Interfaces* **26**(6) 68-78.
- [20] Golkar, J., A. Shekhar, S. Buddhavarapu. 1998. Panama canal simulation model. *Proceedings of the 1998 Winter Simulation Conference*. D. J. Medeiros, E. F. Watson, J. S. Carson and M. S. Manivannan, eds., 1229-1237.
- [21] Ryan, N. K. 1998. The future of maritime facility designs and operations. *Proceedings of the 1998 Winter Simulation Conference*. D. J. Medeiros, E. F. Watson, J. S. Carson and M. S. Manivannan, eds., 1223-1227.
- [22] Merrick, J. R. W., J. R. van Dorp, J. Harrald, T. Mazzuchi, J. Spahn, M. Grabowski. 2000. A systems approach to managing oil transportation risk in Prince William Sound. *Systems Engineering* **3**(3) 128-142.
- [23] Merrick, J., J. R. van Dorp, T. Mazzuchi, J. Harrald, J. Spahn, M. Grabowski. 2002. The Prince William Sound Risk Assessment. *Interfaces* **32**(6) 25-40.
- [24] van Dorp, J. R., J. R. W. Merrick, J. Harrald, T. Mazzuchi, M. Grabowski. 2001. A Risk Management Procedure for the Washington State Ferries. *Risk Analysis* **21**(1) 127-142.
- [25] Banks, J., J. S. Carson, B. L. Nelson, D. M. Nicol. 2000. *Discrete-Event System Simulation*. Prentice Hall, Upper Saddle River, NJ.

- [26] Cooke, R. M. 1991. *Experts in Uncertainty: Expert Opinion and Subjective Probability in Science*. Oxford University Press, Oxford UK.
- [27] Frey, H., D. Burmaster, 1999. Methods for characterizing variability and uncertainty: comparison of bootstrap simulation and likelihood-based approaches. *Risk Analysis* **19**(1) 109-130.
- [28] Apostolakis, G. 1978. Probability and risk assessment: the subjectivist viewpoint and some suggestions. *Nuclear Safety* **19** 305-315.
- [29] Hofer, E. 1996. When to separate uncertainties and when not to separate. *Reliability Engineering and System Safety* **54** 113-118.
- [30] Hora, S. 1996. Aleatory and epistemic uncertainty in probability elicitation with an example from hazardous waste management. *Reliability Engineering and System Safety* **54** 217-223.
- [31] Winkler, R. 1996. Uncertainty in probabilistic risk assessment. *Reliability Engineering and System Safety* **54** 127-132.
- [32] Anderson, E., D. Hattis, N. Matalas, V. Bier, S. Kaplan, D. Burmaster, S. Conrad and S. Ferson. 1999. Foundations. *Risk Analysis* **19**(1) 47-68.
- [33] Berger, J. 1985. *Statistical decision theory and Bayesian analysis*. 2nd edition. New York: Springer-Verlag.
- [34] Bernardo, J., A. Smith. 1994. *Bayesian Theory*. Chichester, UK: Wiley.
- [35] Gilks, W., S. Richardson, D. Spiegelhalter. 1996. *Markov Chain Monte Carlo in practice*. London: Chapman and Hall.

- [36] Glynn, P. 1986. Problems in Bayesian analysis of stochastic simulation. *Proceedings of the Winter Simulation Conference*, ed. J. Wilson, J. Henrikson and S. Roberts, 376-383.
- [37] Chen, H., B. Schmeiser. 1995. Monte Carlo estimation for guaranteed coverage non-normal tolerance intervals. *Journal of Statistical Computation and Simulation* **51**(2-4) 223-238.
- [38] Chen, H. 1996. A lower bound for the correct subset selection probability and its application in discrete event simulations. *IEEE Transactions on Automatic Control* **41**(8) 1227-1231.
- [39] Cheng, R. 1999. Regression metamodeling in simulation using Bayesian methods. *Proceedings of the Winter Simulation Conference*. Ed. P. Farrington, H. Nembhard, D. Sturrock and G. Evans, 330-335.
- [40] Chick, S. 1997. Bayesian analysis for simulation input and output. *Proceedings of the Winter Simulation Conference*. Ed. S. Adradottir, K. Healy, D. Withers and B. Nelson, 253-260.
- [41] Chick, S. 2001. Input Distribution Selection for Simulation Experiments: Accounting for Input Uncertainty. *Operations Research* **49**(5) 744-758.
- [42] Chick, S., K. Inoue. 2001. New two-stage and sequential procedures for selecting the best simulated system. *Operations Research* **49**(5) 732-743.
- [43] Barton, R. R., L. W. Schruben 2001. Resampling Methods for Input Modeling. *Proceedings of the Winter Simulation Conference*. Eds. B. A. Peters, J. S. Smith, D. J. Medeiros and M. W. Rohrer, 372-378.

- [44] Chick, S. 2000. Bayesian methods for simulation. *Proceedings of the Winter Simulation Conference*. Ed. J. Joines, R. Barton, K. Kang & P. Fishwick, 109-118.
- [45] Merrick, J., J. R. van Dorp, J. P. Blackford, G. L. Shaw, J. Harrald, T. Mazzuchi. 2003. A traffic density analysis of proposed ferry service expansion in San Francisco Bay using a maritime simulation model. *Reliability Engineering and System Safety* **81**(2) 119-132.
- [46] Law, A., W. Kelton 2001. *Simulation Modeling and Analysis*. 3rd edition. New York: McGraw Hill.
- [47] Kass, R., A. Raftery. 1995. Bayes Factors. *Journal of the American Statistical Association* **90**(430) 773-795.
- [48] Gelfand, A. 1996. Model determination using sampling-based methods. *Markov Chain Monte Carlo in Practice*. Eds. W. Gilks, S. Richardson and D. J. Spiegelhalter, London: Chapman and Hall, 145-161.
- [49] Spiegelhalter, D. J., N. G. Best, B. P. Carlin, A. van der Linde. 2002. Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B* **64**(4) 583-639.
- [50] Spiegelhalter, D., A. Thomas, N. Best, W. Gilks. 1996. *BUGS: Bayesian inference Using Gibbs Sampling, Version 0.5*. MRC Biostatistics Unit, Cambridge, UK.
- [51] Ferguson, T. S. 1974. Prior distributions on spaces of probability measures. *Annals of Statistics* **2** 615-629.

Table I. The Decision Information Criteria values for the chosen arrival process.

	\bar{D}	p_D	DIC
exponential	877.778	1.010	878.787
Weibull	838.202	1.750	839.952
gamma	837.349	2.025	839.374
log normal	847.737	1.999	849.736

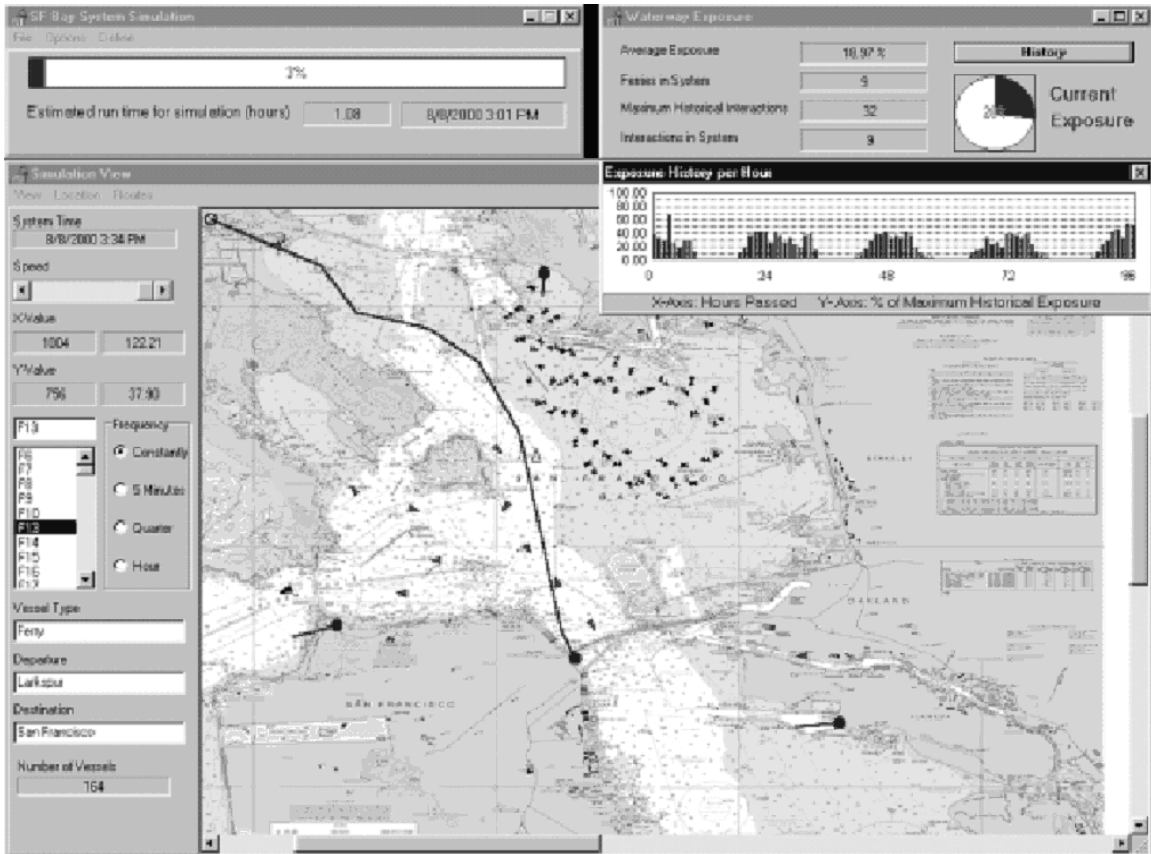


Figure 1. The simulation of San Francisco Bay.

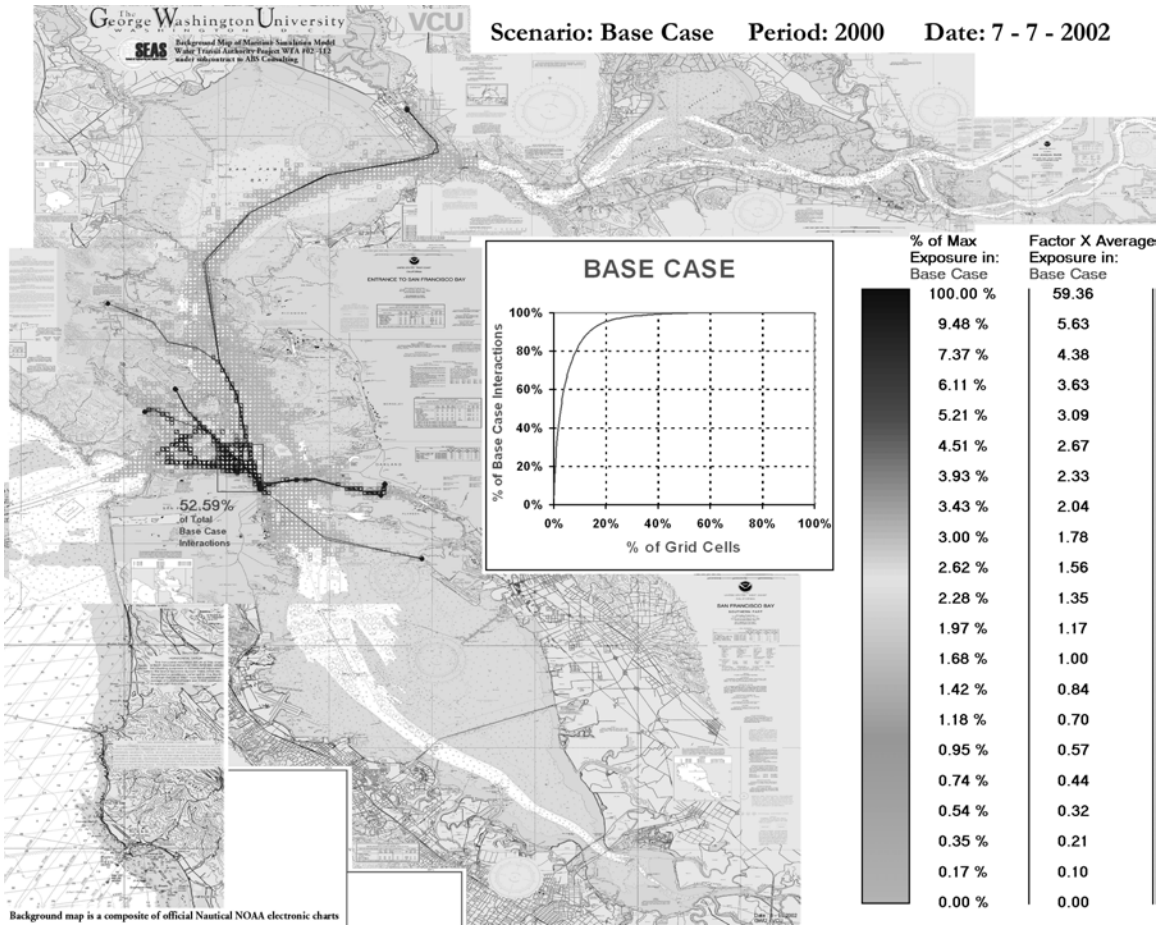


Figure 2. The simulation results for the current ferry system.

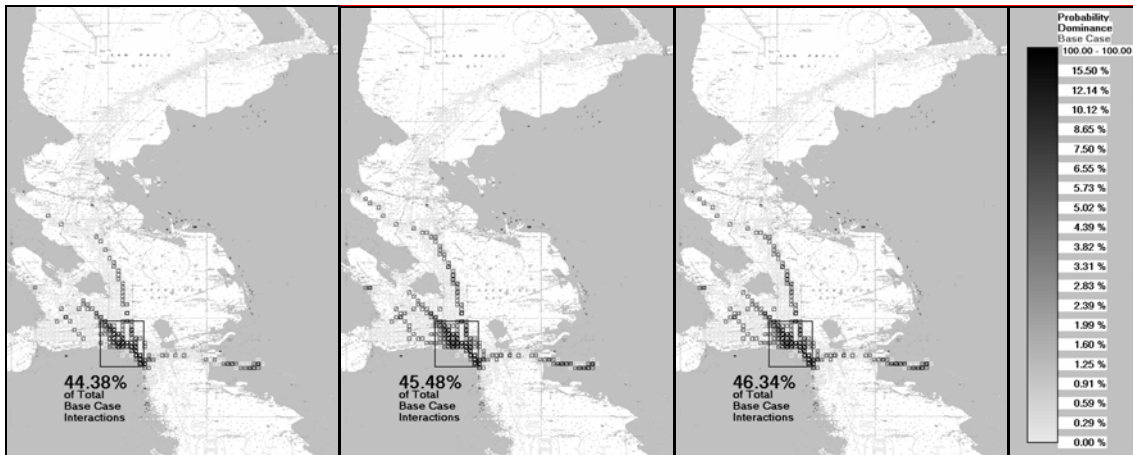


Figure 3. 5th, 50th and 95th Percentiles of the Daily Expected Number of Interactions in the Base Case.

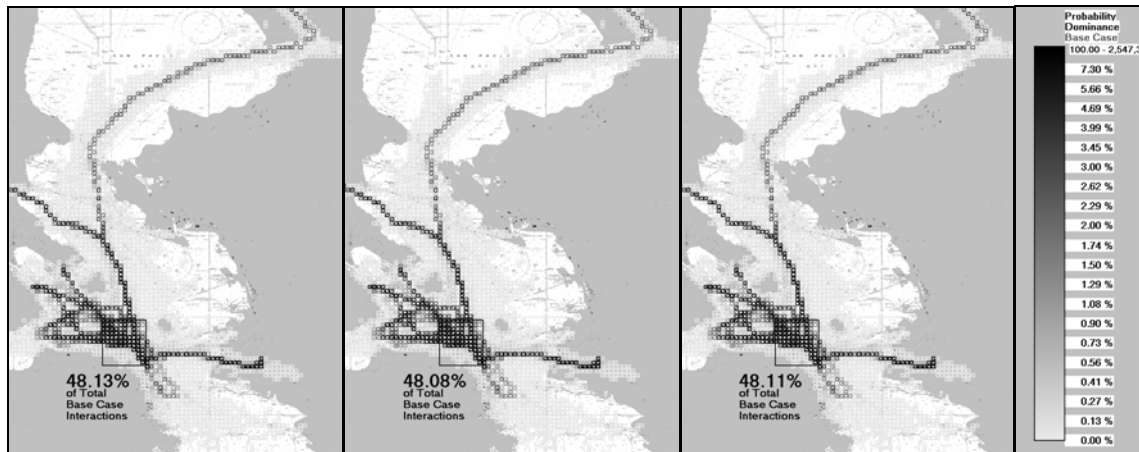


Figure 4. 5th, 50th and 95th Percentile Maps of the Yearly Expected Number of Interactions in the Base Case

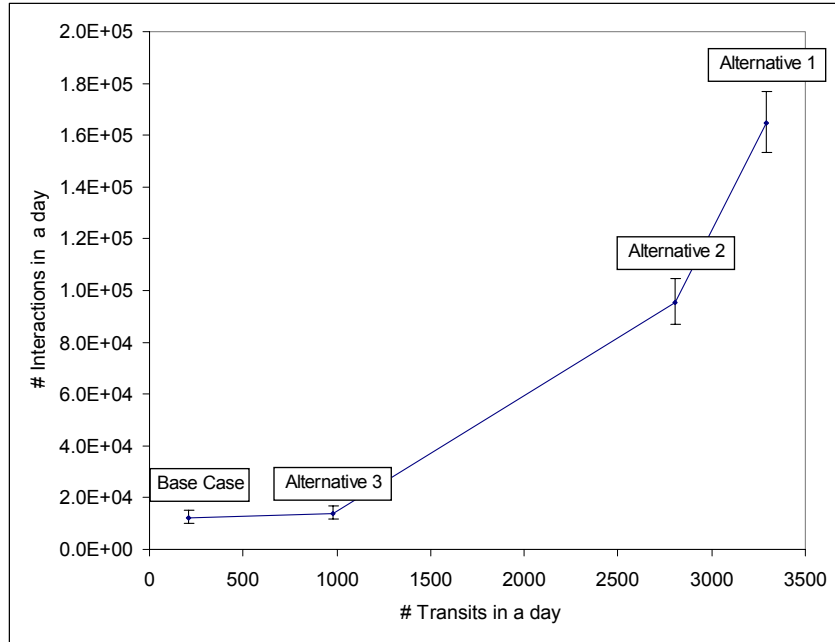


Figure 5. Daily Expected Number of Interactions for the Four Scenarios with 90% prediction intervals.

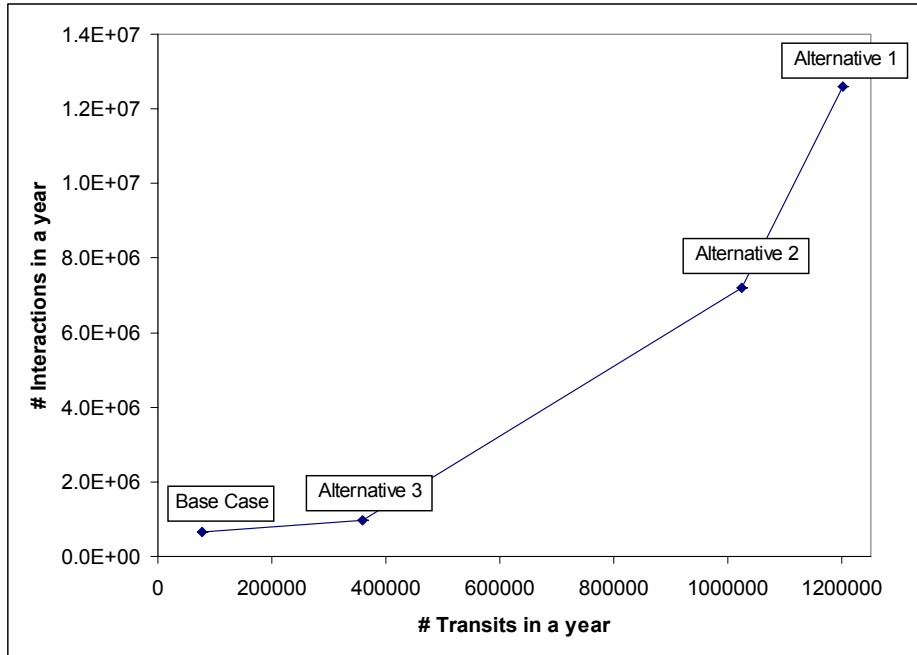


Figure 6. Yearly Expected Number of Interactions for the Four Scenarios with 90% prediction intervals (not visible).

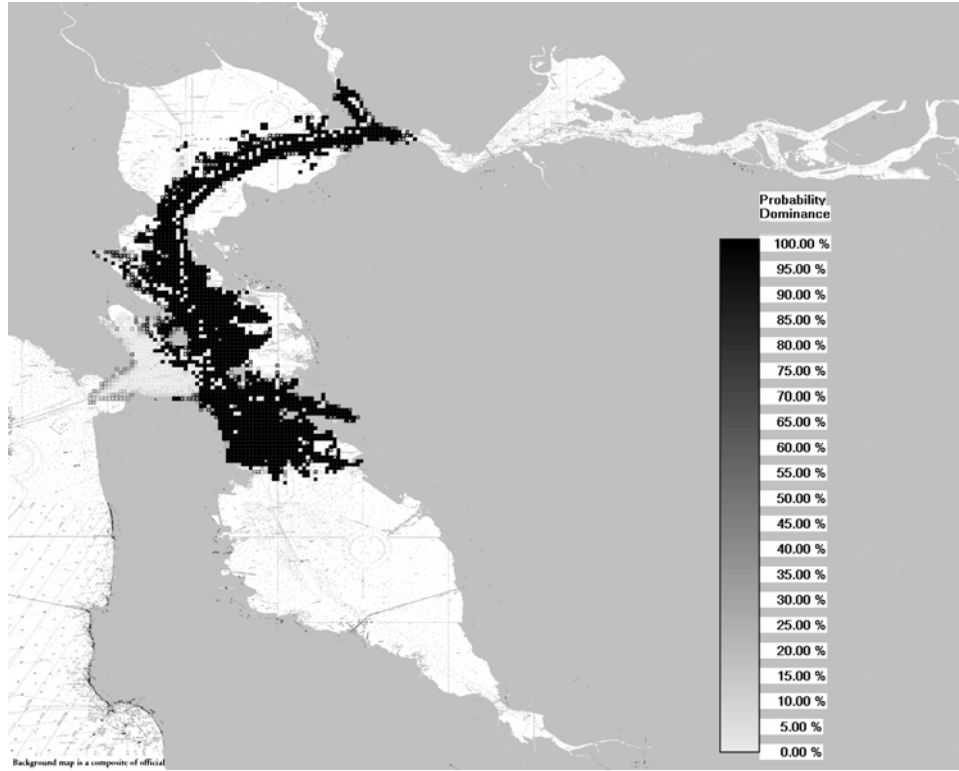


Figure 7. Base Case compared to Alternative 3




Figure 8. Alternative 2 compared to Alternative 1