Paper No. 03-4021

Systematic Investigation of Variability due to Random Simulation Error in an Activity-Based Microsimulation Forecasting Model

Joe Castiglione, Joel Freedman, and Mark Bradley

A key difference between stochastic microsimulation models and more traditional forms of travel demand forecasting models is that microsimulation-based forecasts change each time the sequence of random numbers used to simulate choices is varied. To address practitioners' concerns about this variation, a common approach is to run the microsimulation model several times and average the results. The question then becomes: What is the minimum number of runs required to reach a true average state for a given set of model results? This issue was investigated by means of a systematic experiment with the San Francisco model, a microsimulation model system used in actual planning applications since 2000. The system contains models of vehicle availability, day pattern choice, tour time-of-day choice, destination choice, and mode choice. To investigate the variability of the forecasts of this system due to random simulation error, the model system was run 100 times, each time changing only the sequence of random numbers used to simulate individual choices from the logit model probabilities. The extent of random variability in the model results is reported as a function of two factors: (*a*) the type of model (vehicle availability, tour generation, destination choice, or mode choice); and (*b*) the level of geographic detail—transit at the analysis zone level, neighborhood level, or countywide level. For each combination of these factors, it is shown graphically how quickly the mean values of key output variables converge toward a stable value as the number of simulation runs increases.

A key difference between stochastic microsimulation models and more traditional forms of travel demand forecasting models is that, in microsimulation models, the forecasts change each time the sequence of random numbers used to simulate choices is varied. This can be a disturbing feature for practitioners, who are used to having the forecast results vary only if model inputs are changed. Instead of running the model several times and reporting the range of potential outcomes for policy analysis, a common approach is to either fix the random number seeds at some arbitrary value and report only one potential state or set of model results, or run the microsimulation model several times and average the results from the runs. The question then becomes: How many times does the model system need to be run to result in confidence that the average results will be stable and representative of a true average? This question is highly relevant to urban planning practitioners seeking to use stochastic microsimulation models and will become increasingly important as these models are developed and applied in other regions. The limited research on this subject by Veldhuisen et al. (*1*) finds a negligible impact of Monte Carlo error on the variance of large-scale microsimulation model results. This issue is more thoroughly investigated here by means of a systematic experiment with a microsimulation model system that has been used in actual planning applications over the last 2 years.

J. Castiglione, San Francisco County Transportation Authority, 100 Van Ness Avenue, 25th Floor, San Francisco, CA 94102. J. Freedman, Systems Analysis Group, PBConsult, 400 Southwest Sixth Avenue, Suite 802, Portland, OR 97204. M. Bradley, Mark Bradley Research & Consulting, 129 Natoma Avenue, Apartment C, Santa Barbara, CA 93101.

SAN FRANCISCO MODEL AND MICROSIMULATION FRAMEWORK

One of the first activity-based microsimulation models used extensively in planning is the model system created by Cambridge Systematics and Parsons Brinckerhoff for the San Francisco County Transportation Authority, completed in 2000. The San Francisco model was developed to accurately represent the complexity of the destination, temporal, and modal options in San Francisco and to provide detailed information about travelers making discrete choices. These objectives led to the development of an activity-based model that uses synthesized population as the basis for decision making instead of zonal-level aggregate data sources. The model system uses the full-day pattern activity modeling approach introduced by Bowman and Ben-Akiva (*2*). The main feature of the full-day pattern approach is that it simultaneously predicts the main components of all of a person's travel across the day.

In the San Francisco model, a microsimulation framework is applied to individuals and households making vehicle ownership, trip pattern (*3*), and trip destination and mode choices (*4*); many of these models are logit formulations. A Monte Carlo method is used to select outcomes according to these logit model probabilities based on random number draws. Each time the sequence of random numbers used to simulate choices is varied, the model result, or end state of the model, may change. Each model component contains a pseudo–random number generator that uses an externally provided random seed. In this investigation, the external seeds were varied explicitly, which causes a different sequence of random numbers to be generated. The random number generator used in the San Francisco County Transportation Authority models is distributed by Borland and has a period of 2^{32}.

Figure 1 presents the San Francisco model structure. The model system predicts the choices for a full, representative sample of residents of San Francisco County, almost 800,000 simulated individual person-days of travel. The following sections briefly describe the primary components of the San Francisco model and how random numbers affect outcomes.

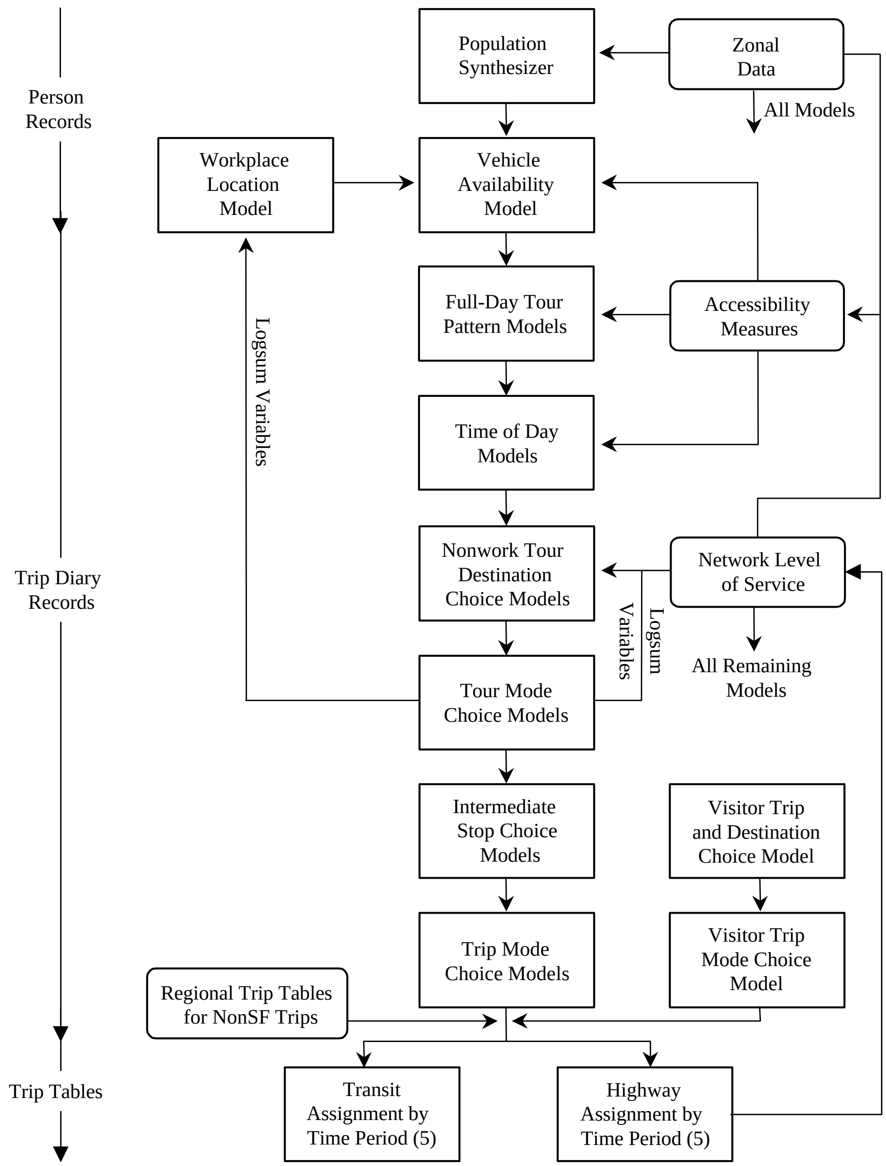

FIGURE 1 San Francisco (SF) model components (*3*). (5) refers to five time periods—early a.m., a.m., midday, p.m., and evening.

SYNTHETIC SAMPLE GENERATOR

This program generates a full synthetic population for a base year or forecast year. Households are categorized simultaneously by household size/workers (nine groups), age of head of household (three groups), and income class (four groups). In this investigation, the synthetic sample is held constant.

WORKPLACE LOCATION MODEL

For each worker in the synthetic sample, a workplace zone is drawn from 40 sampled zones according to multinomial logit model probabilities based on worker characteristics, mode choice accessibility logsums, and work zone attributes. Random numbers, used in the selection of the sample set of zones and in the selection of the workplace zone, are varied in this investigation.

VEHICLE AVAILABILITY MODEL

For each synthetic household, 0, 1, 2, or 3+ automobiles are chosen according to multinomial logit probabilities based on the characteristics of the household. The random number seed used to control selection of the number of automobiles is varied in this investigation.

FULL-DAY TOUR AND TRIP PATTERN–TIME-OF-DAY MODELS

For each synthetic person, the probability of each full-day pattern (composed of tours and trips), including no travel, is predicted for each person. A random Monte Carlo procedure is used to select a single pattern. This investigation presents the results of varying the seed of this simulation.

TOUR PRIMARY DESTINATION AND MODE CHOICE MODELS

For each tour, a primary destination is chosen from 40 sampled zones according to logit probabilities based on characteristics of the person and tour, zone size, and mode choice logsums; a tour mode is chosen from a set of six potential modes according to logit probabilities based on person and tour characteristics and levels of service between the tour origin and destination. The random numbers varied in this investigation are used in the selection of the sample set of destinations, the selection of the chosen destination, and the chosen mode.

INTERMEDIATE STOP LOCATION CHOICE

For each intermediate stop on each tour, an intermediate stop zone is chosen from 40 sampled zones according to multinomial logit probabilities based on characteristics of the person and tour and the zone size as well as the additional cost of travel between the tour origin and destination imposed by the sampled stop. Random numbers are used to control the selection of the sampled zones and the selection of an intermediate stop. The results of varying the random seeds of this model are not presented here.

TRIP MODE CHOICE

For each tour, a mode is chosen from 11 possible modes according to logit probabilities based on characteristics of the person and tour and levels of service between the trip origin and destination zone conditioned by the chosen tour mode. A random number is used to select the chosen mode. The results of varying the random seeds of this model are not presented here.

METHODOLOGY

To investigate the variability of the forecasts of this system due to random simulation error, the model system was run 100 times, changing the sequence of random numbers in each run. The logit probabilities of each component in the model chain varies for each run, because the choice models are conditional on choices made further up the decision chain. For example, tour mode choice and destination choice probabilities are conditional on the predicted tour purpose, number of intermediate stops on the tour, and times of day when the trips in the tour are carried out. For this reason, it would be very difficult to calculate the exact probability distribution for any particular model outcome. In this paper, the extent of random variability in the model results is explored primarily as a function of two factors:

1. Type of model—vehicle availability, tour generation, destination choice, and mode choice (described previously); and
2. Level of geographic detail—zone [transit at the analysis zone (TAZ)] level, neighborhood level, and countywide level.

Figure 2 indicates the levels of geographic detail used in this analysis. There are 766 TAZs in San Francisco and 1,740 in the region. There are also 26 neighborhoods, within which the TAZs nest, and 1 county, within which the neighborhoods nest. The TAZ and neighborhood results presented in the graphs and table that follow refer to the two zones and two neighborhoods highlighted in black and in gray in Figure 2. Table 1 presents descriptive statistics related to all graphs and indicates which selected TAZ or neighborhood is represented in a given graph. This table is included because the graphs present only the percent difference between the sample mean and the final mean and not the actual values. This output format was chosen to facilitate comparisons across models and geographic scales.

For combinations of model type and geography, it is shown graphically how quickly the mean values of key output variables converge toward a stable value as the number of simulation runs increases. On some graphs, it is also shown how quickly the standard errors of the mean values converge based on the distribution of simulation results up to that point. Although this is a fairly simple analysis, it is a clear way of showing how the confidence in the mean results varies with the number of simulation runs for this model system. The authors have attempted to present statistics that typically would be of interest to practitioners using microsimulation models. While the aggregation of individual choices, even at a TAZ level, necessarily produces an average statistic, it was believed that it would be of little use to practitioners to present disaggregate individual choice results.

The following formula was used to calculate the sample means after n runs:

$$\overline{X}_n = \frac{\sum_{i=1}^{n} X}{n} \tag{1}$$

where X represents the value for a given run and n represents the run number.

The following equation was used to calculate the percent difference between the sample means and the final mean:

$$\%\text{diff} = \frac{\overline{X}_n - \overline{X}_{100}}{\overline{X}_{100}} \times 100 \tag{2}$$

where $\overline{X}_n$ represents the sample mean after n runs, and $\overline{X}_{100}$ represents the sample mean after 100 runs.

The following equation was used to calculate the standard error of the sample means:

$$\sigma_{M_n} = \frac{\sqrt{\frac{\sum_{i=1}^{n} (X_i - \overline{X}_n)^2}{n}}}{\sqrt{n}} \tag{3}$$

where n represents the run number, and $\overline{X}_n$ represents the sample mean after n runs.

The following equation was used to calculate the percent difference between the standard error of the sample mean and the standard error of the final mean:

$$\%\text{diff } \sigma_M = \frac{\sigma_{M_n} - \sigma_{M_{100}}}{\sigma_{M_{100}}} \times 100 \tag{4}$$

where σ_{M_n} represents the standard error of the mean after n runs, and σ_{M100} represents the standard error of the mean after 100 runs.

RESULTS

The key question investigated is not whether the model components will converge to a stable mean but how quickly the models converge to this mean, acknowledging that central limit theorem indicates that

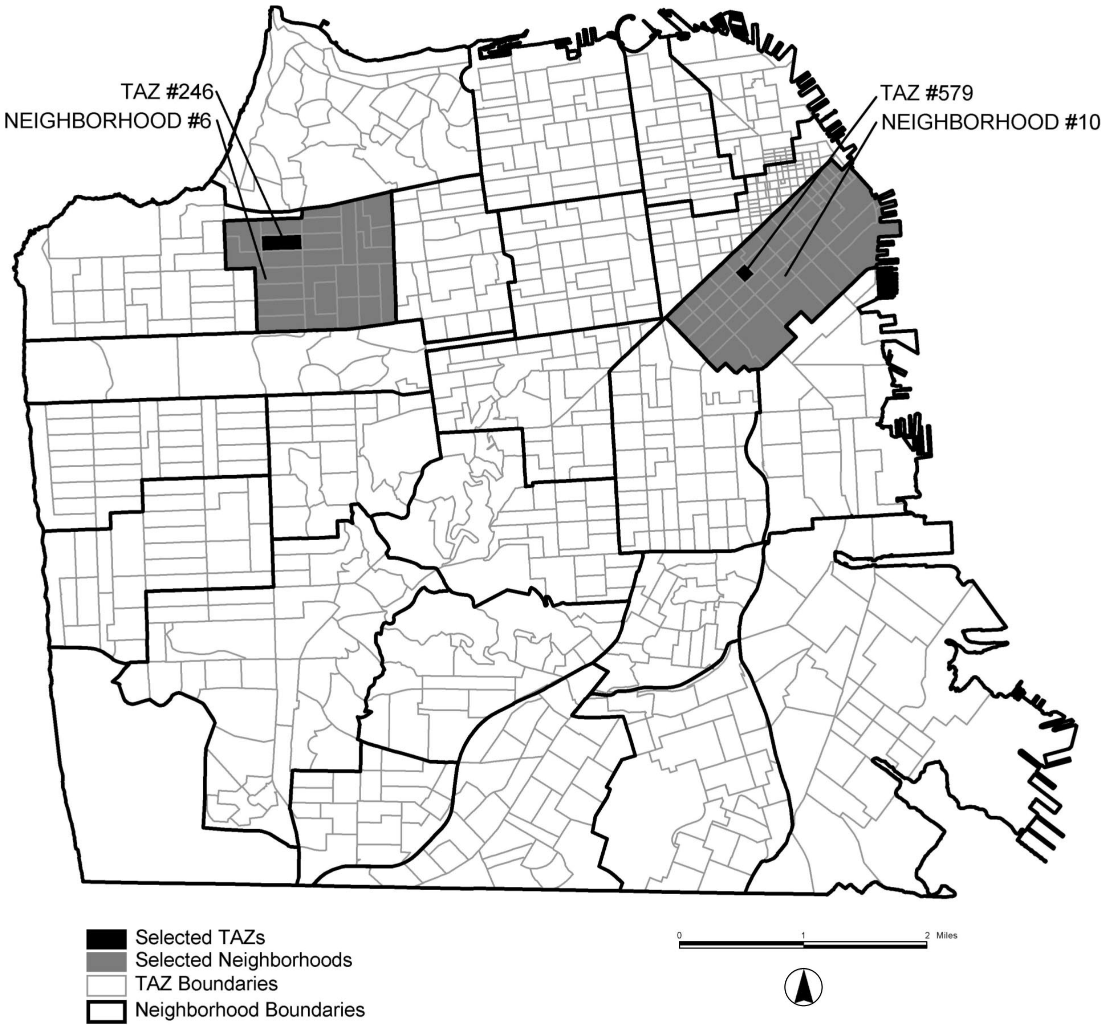

FIGURE 2 San Francisco model TAZ and neighborhood map.

the deviation from the mean will decrease as the number of observations increases. In addition, a comparison of the convergence at different geographic scales was sought, with an understanding that higher levels of geographic aggregation will necessarily result in greater stability. However, a Monte Carlo realization of a behavioral microsimulation model can reveal useful information about the stability of complex interactive models.

As expected, the results at the countywide level are much more stable than those at the individual zone level. The variability at the neighborhood level (each neighborhood contains about 30 TAZs) falls in between the two but is generally much more stable than the individual zone level.

Also not surprisingly, the results for the models with many alternatives, such as destination choice, show more random variability than do the results for models with fewer alternatives, such as vehicle availability and tour generation. For destination and mode choice, the results are more stable when looking at trips from a particular origin TAZ or to a particular destination TAZ than when looking at trips between a particular origin–destination pair. Even for this most extreme test, however, the forecasts become fairly stable once one moves from the TAZ-to-TAZ level of detail out to the neighborhood-to-neighborhood level.

In general, the more rare a particular choice or outcome is, the less stability will be observed in the results. As a result, there tends to be higher variability in transit mode shares than in automobile mode shares in areas where transit mode shares are low. The same applies for the percentage of zero-vehicle households in areas where almost all households own cars. In San Francisco, the shares of transit use and zero-vehicle households are both higher than in most areas, so the forecasts of those shares should tend to be more stable than in a typical city. Also, cities and regions that are much larger (more TAZs) and with more complex land use patterns tend to show more random variability and thus require either more simulation runs or a somewhat higher spatial aggregation before achieving stable results.

These findings have distinct implications for how these models are applied. While travel demand model outputs are often analyzed at aggregate geographic levels, there is increasing interest in relying on model outputs at ever-finer geographic levels. In the microsimulation context, it is critical to understand how many runs are necessary to ensure confidence at these different geographic scales. In general, the more specific the policy or statistic of interest, the smaller the geographic level, or the rarer a particular choice of interest, the greater the number of runs required.

TABLE 1 Descriptive Statistics for Figures 3 to 13

		MEAN	MIN	MAX	STDEV	STDERR
FIGURE 3 % 0-Vehicle Households	County	26.41%	26.29%	26.63%	0.06%	0.01%
	Neigh (#6)	17.02%	16.29%	17.81%	0.26%	0.03%
	TAZ (#246)	15.23%	11.34%	18.35%	1.37%	0.14%
FIGURE 4 Trips per Person	County	3.96	3.95	3.97	0.004	0.000
	Neigh (#6)	4.05	4.01	4.11	0.004	0.000
	TAZ (#246)	4.14	3.82	4.56	0.106	0.011
FIGURE 5 Neighborhood Tour Destinations (All Purposes)	Neigh (#10)	96,669	96,088	97,643	316.26	31.63
FIGURE 6 Neighborhood-Neighborhood Tours (All Purposes)	Neigh (#6) to Neigh (#10)	3,427	3,305	3,590	57.30	5.73
FIGURE 7 Neighborhood-Neighborhood Tours (Work Purpose)	Neigh (#6) to Neigh (#10)	1,811	1,724	1,907	40.19	4.02
FIGURE 8 Auto Tours (by Origin)	County	633,150	630,817	635,552	948.50	94.85
	Neigh (#6)	30,991	30,502	31,403	186.86	18.69
	TAZ (#246)	1,071	981	1,211	33.48	3.35
FIGURE 9 Transit Tours (by Origin)	County	211,538	210,616	212,570	438.28	43.83
	Neigh (#6)	9,572	9,386	9,814	97.54	9.75
	TAZ (#246)	312	276	362	17.32	1.73
FIGURE 10 Nonmotorized Tours (by Origin)	County	243,775	242,371	245,771	571.24	57.12
	Neigh (#6)	9,958	9,670	10,194	104.88	10.49
	TAZ (#246)	308	253	358	20.36	2.04
FIGURE 11 Tours by Mode (by origin=TAZ #246)	Auto	1,071	981	1,211	33.48	3.35
	Transit	312	276	362	17.32	1.73
	Nonmotorized	308	253	358	20.36	2.04
FIGURE 12 Tours by Mode (by destination=Neighborhood #10)	Auto	38,332	37,795	38,830	192.45	19.24
	Transit	33,907	33,197	34,391	187.23	18.72
	Nonmotorized	24,430	24,065	24,832	162.42	16.24
FIGURE 13 Tours by Mode (by destination=TAZ #579)	Auto	565	490	624	22.66	2.27
	Transit	362	326	408	17.30	1.73
	Nonmotorized	410	336	455	21.57	2.16

NOTE: This table reflects actual values for each statistic of interest from the 100 runs and is included because the graphs present only the percent difference between the sample mean and the final mean. STDEV = standard deviation; STDERR = standard error; Neigh = neighborhood.

The following sections examine the stability of four of the eight San Francisco model components: vehicle availability, tour generation, tour destination choice, and tour mode choice. A final section examines stability when the 100-run sample is divided into four smaller 25-run samples. Figures 3 to 13 are in the following format: along the *x*-axis is the run number completed. A total of 100 runs of each model component were performed, and the results are presented sequentially. On the *y*-axis is the percentage difference between the mean value after a particular run and the final mean value from all 100 runs. The TAZ and neighborhood values presented in the following figures are for a single set of typical subcounty origin and destinations, indicated in Figure 2.

VEHICLE AVAILABILITY

For the vehicle availability model, one primary measure was summarized at the three geographic levels: the percentage of households with zero vehicles. This measure proved to be relatively stable at all three geographic levels.

Figure 3 indicates the percent difference from the final mean for each run for the percentage of zero-vehicle households. At both the county level and neighborhood level, the sample mean is never more than 0.5% different than the final mean after 100 runs. The TAZ level indicates much more variation but still well within reason—after only 10 runs, the mean is within 2.5% of the final result after 100 runs.

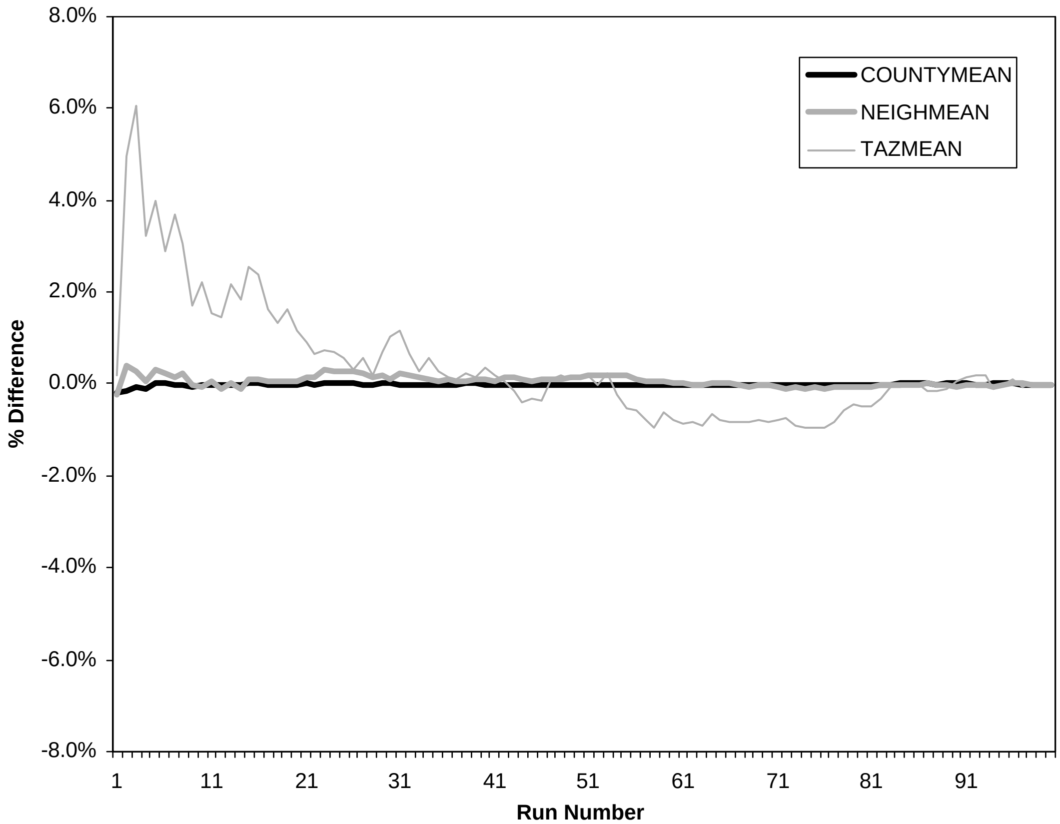

FIGURE 3 Percent zero-vehicle households (all levels), percent difference from final mean.

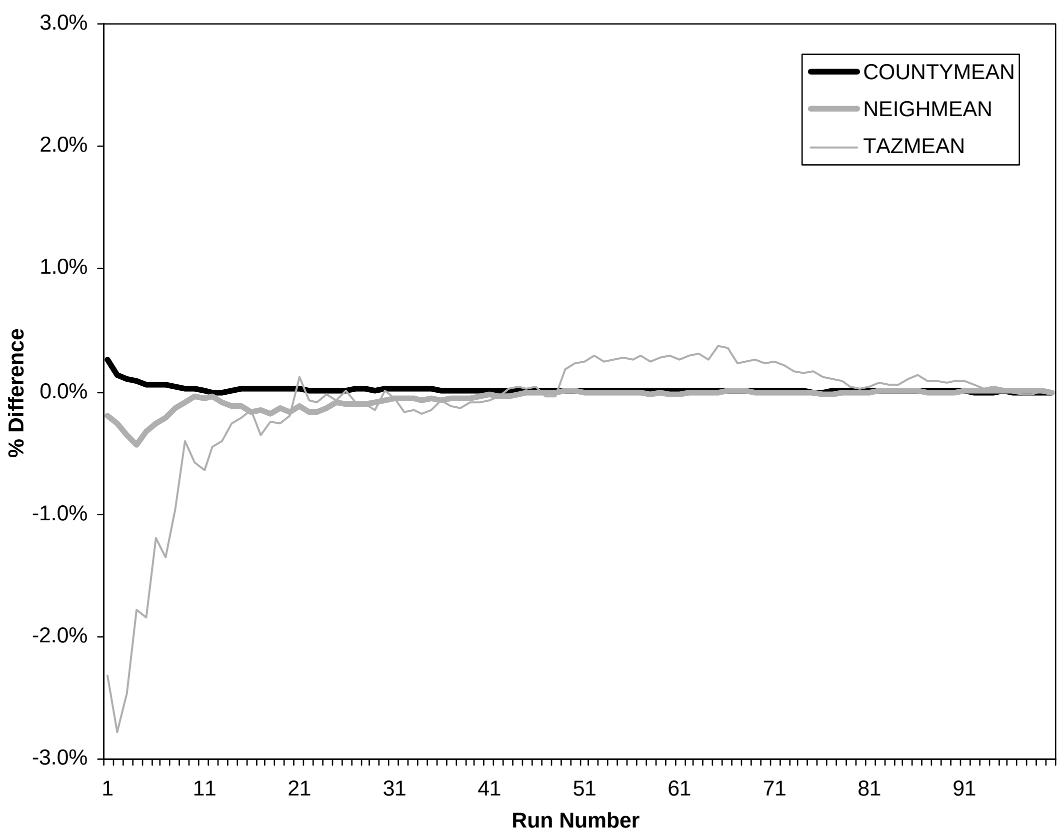

FIGURE 4 Trips per person (all levels), percent difference from final mean.

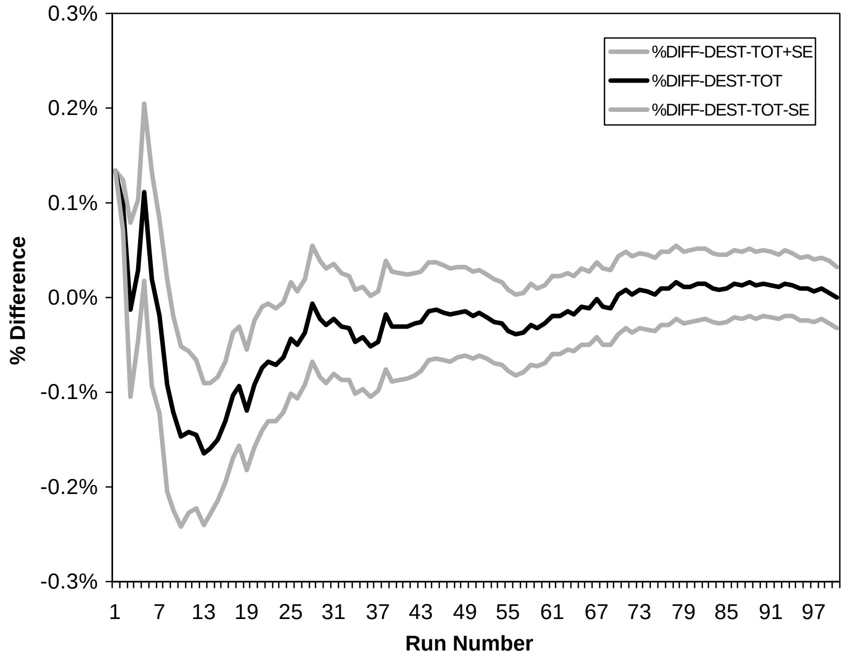

FIGURE 5 Neighborhood tour destinations (all purposes), percent difference from final mean (SE = standard error).

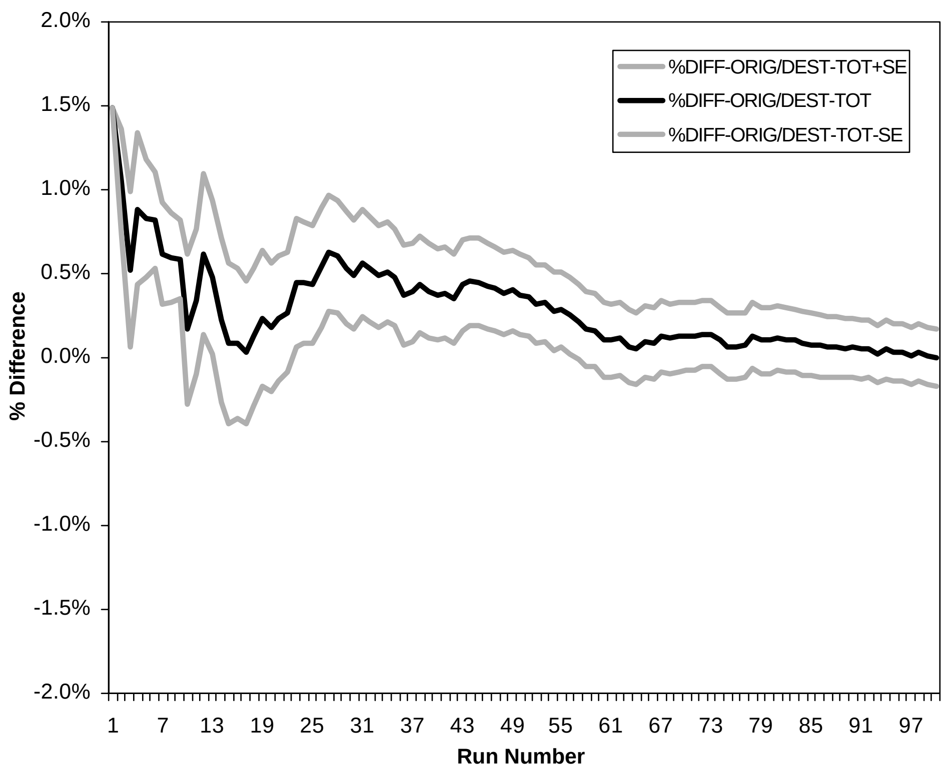

FIGURE 6 Neighborhood-to-neighborhood tours (all purposes), percent difference from final mean.

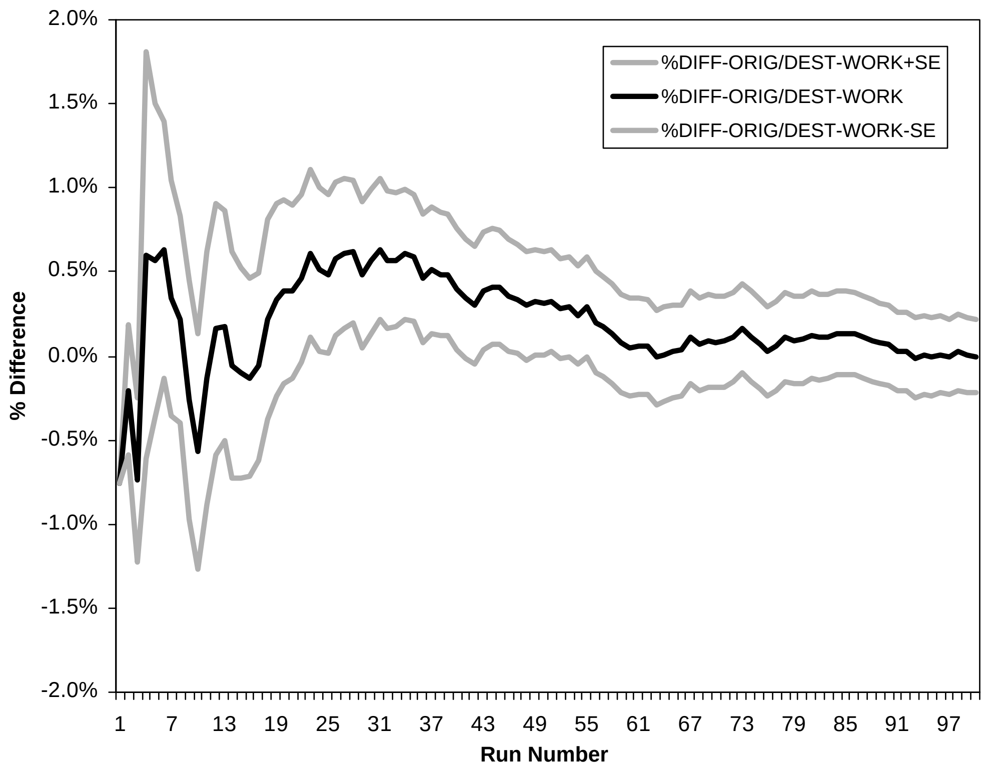

FIGURE 7 Neighborhood-to-neighborhood tours (work purpose), percent difference from final mean.

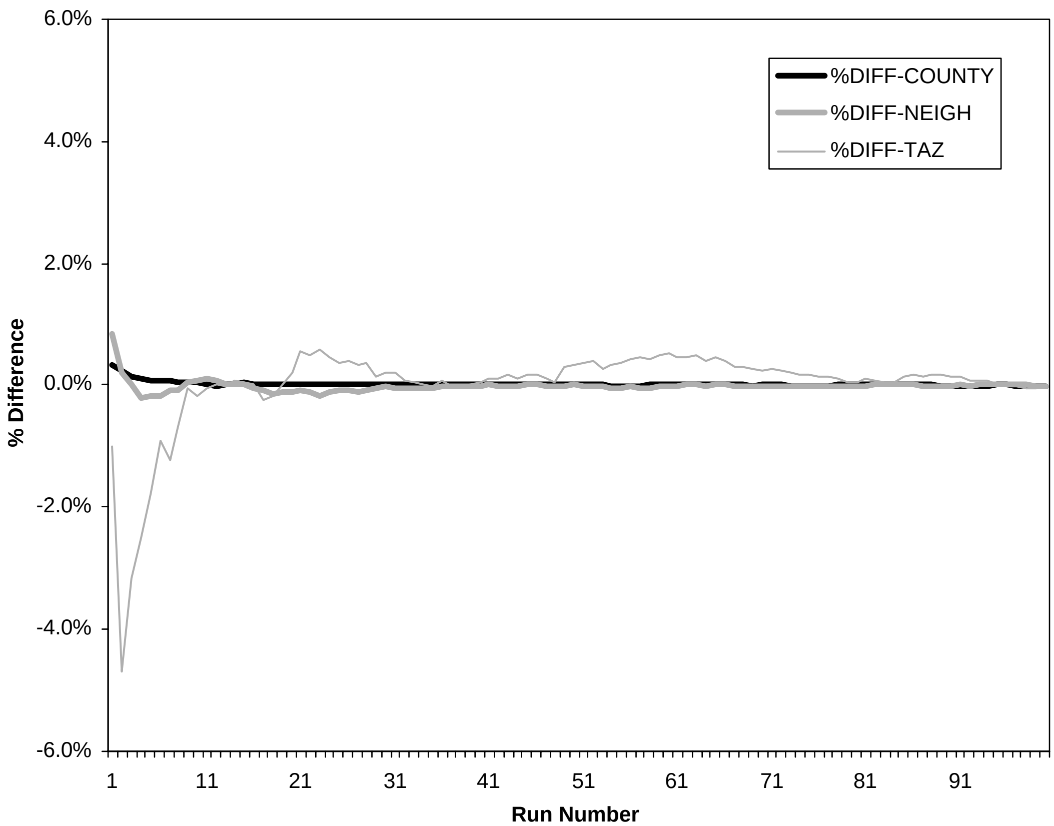

FIGURE 8 Automobile tours (all levels), percent difference from final mean.

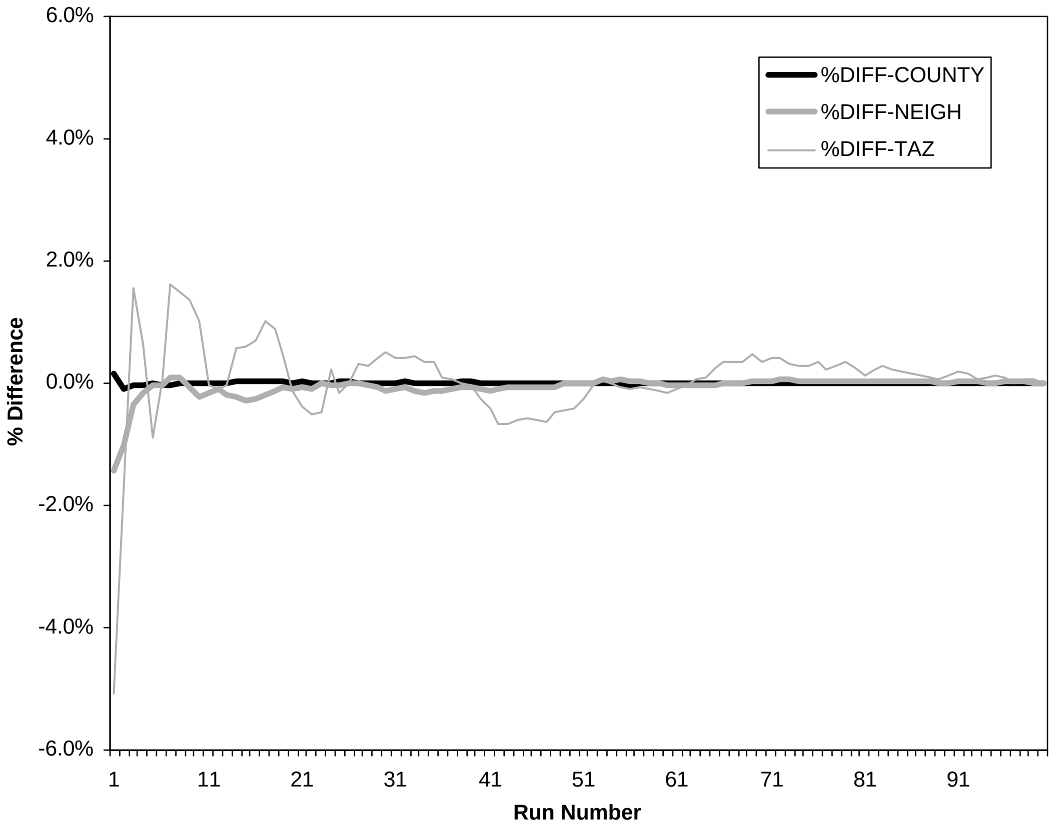

FIGURE 9 Transit tours (all levels), percent difference from final mean.

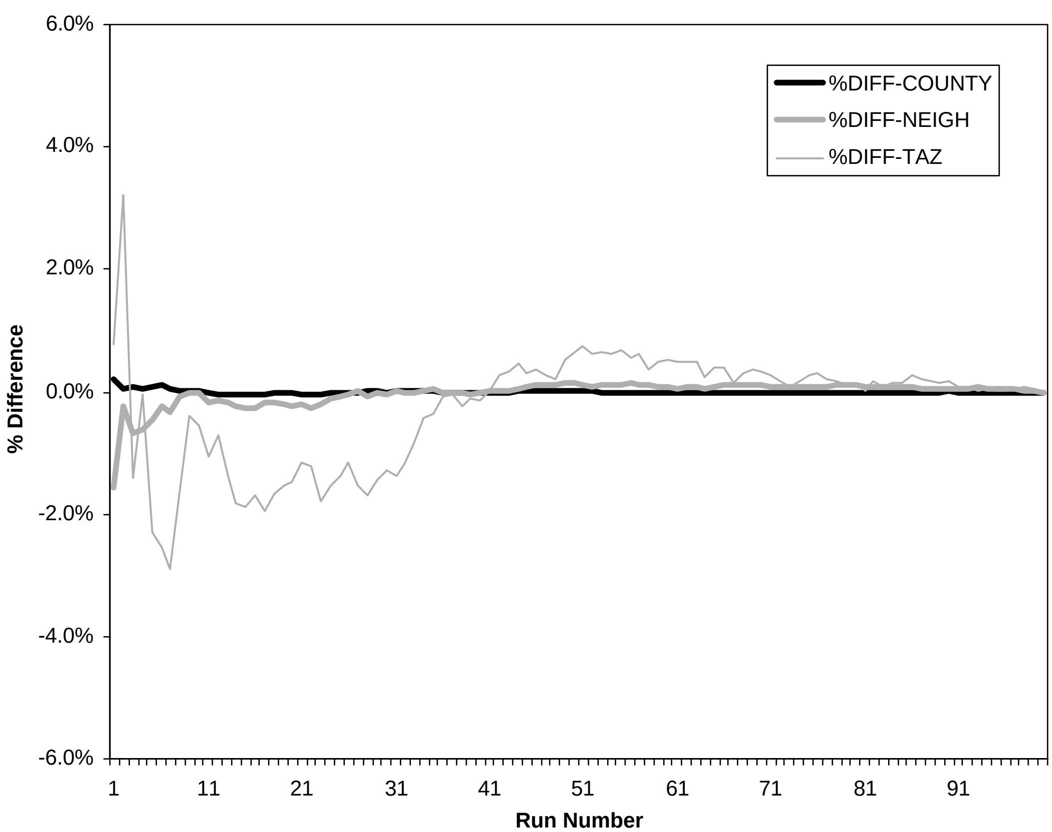

FIGURE 10 Nonmotorized tours (all levels), percent difference from final mean.

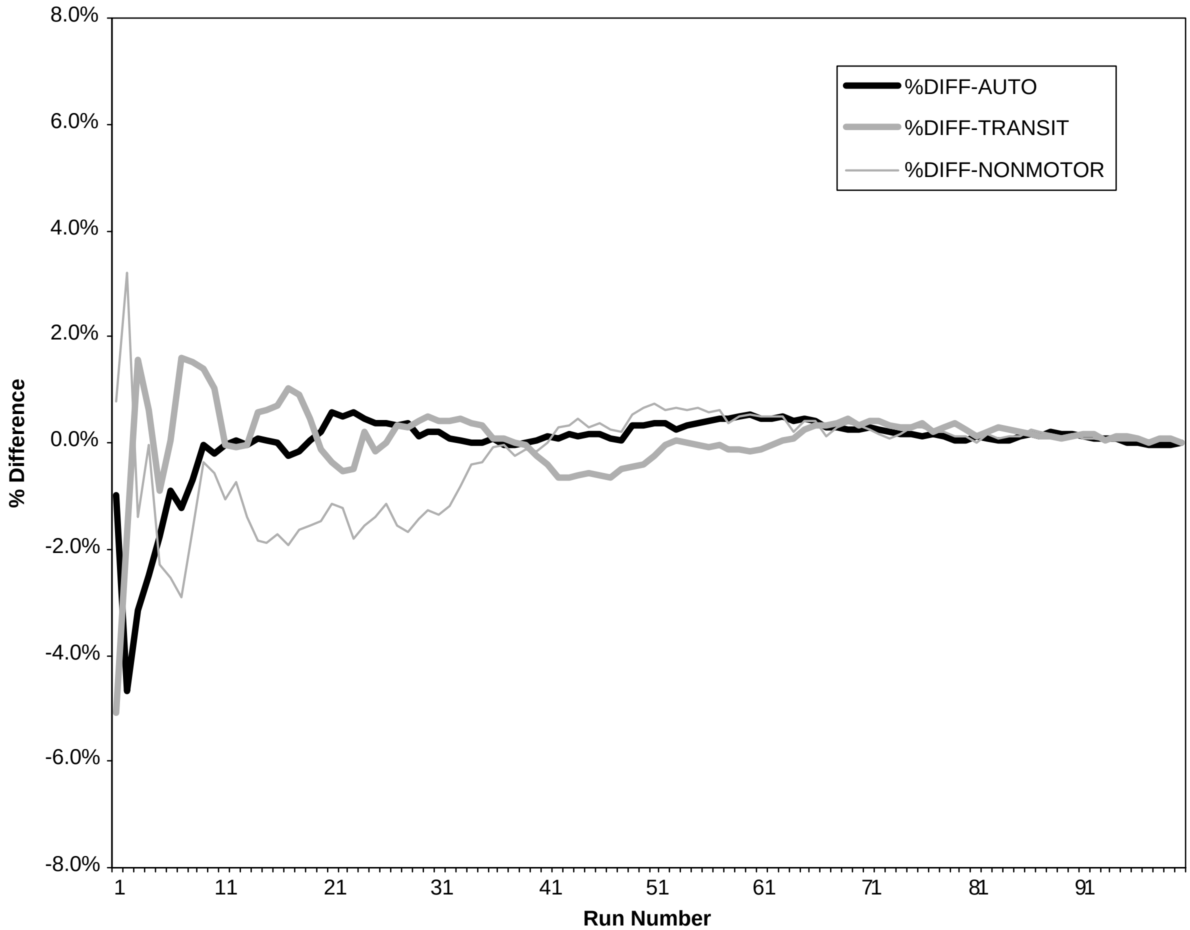

FIGURE 11 Tours by mode (TAZ = origin), percent difference from final mean.

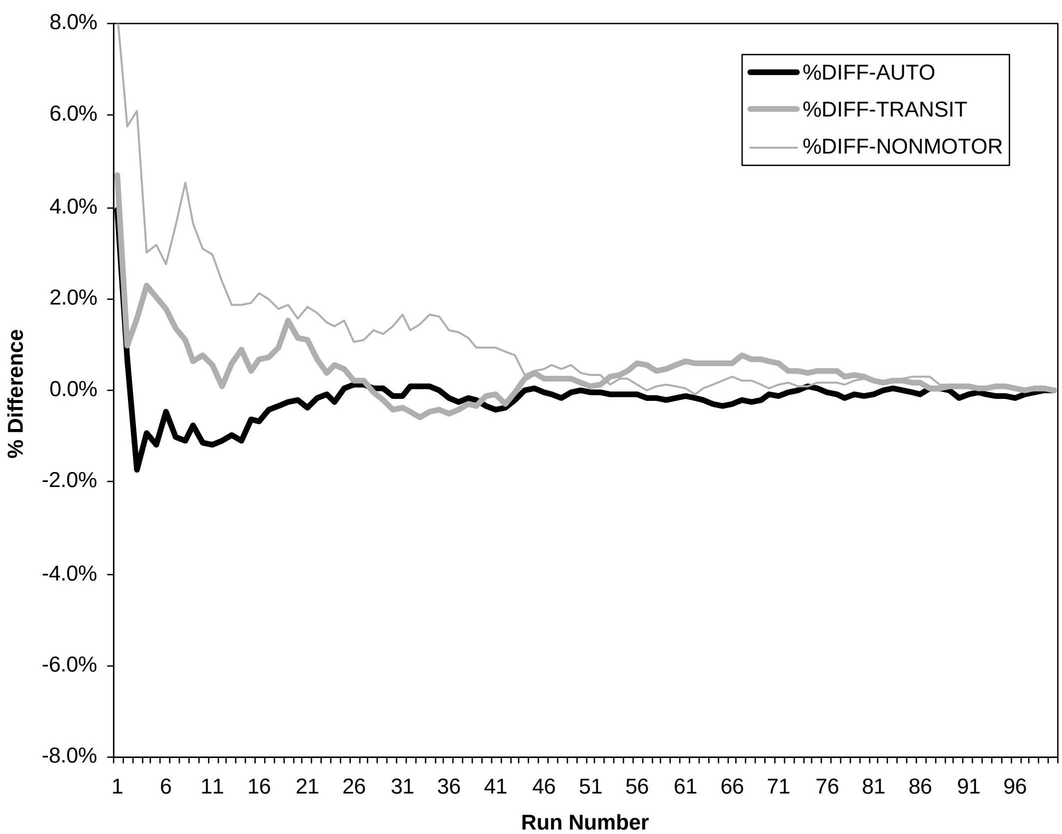

FIGURE 12 Tours by mode (TAZ = destination), percent difference from final mean.

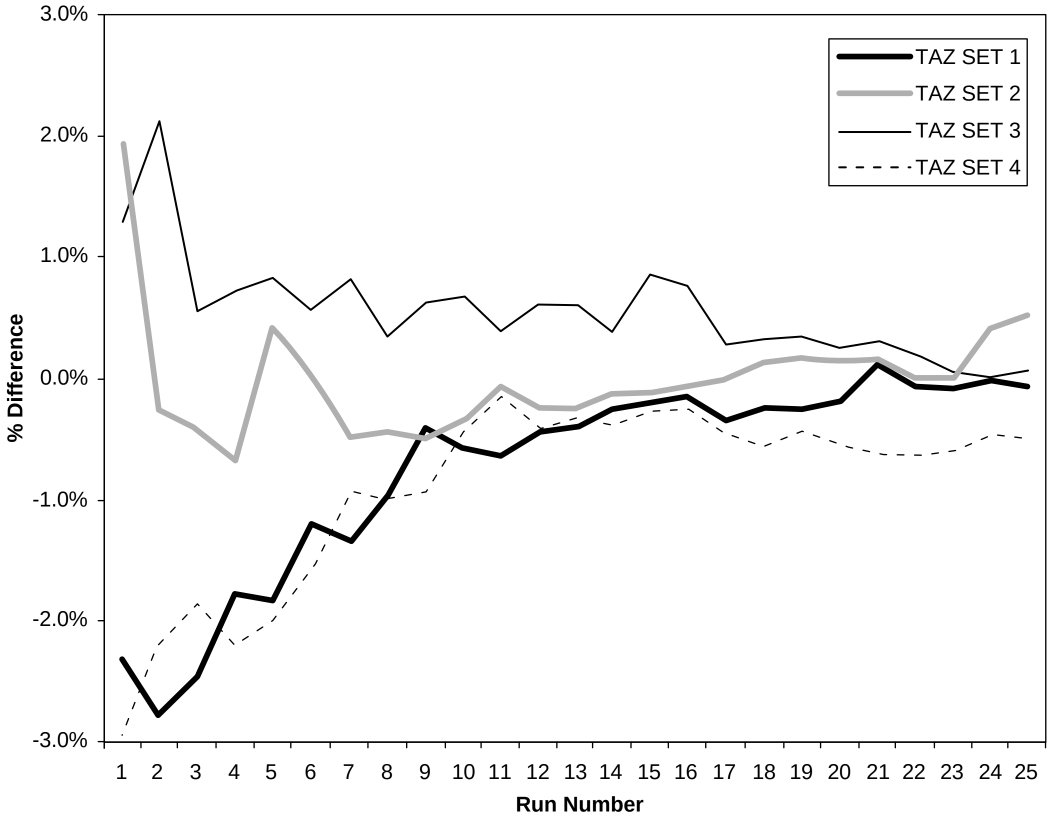

FIGURE 13 Trips per person (TAZ level, four sets), percent difference from final 100-run mean.

TOUR GENERATION

The number of trips per person was used as the primary measure of the stability of the tour generation model. The tour generation model is applied at the level of the individual, and the results presented include trips for all purposes. While the San Francisco model predicts tours as well as individual trips, the pattern for both tours and trips at all geographic levels was the same. Figure 4 indicates the changes in trips per person. As with the vehicle availability model, the results appear relatively stable at all geographic summary levels. However, more variation was observed at the TAZ level than in the vehicle availability model. At the neighborhood and county level, the results never vary more than 0.5% from the result after 100 runs. At the TAZ level, the initial runs are 2% to 3% different than the result after 100 runs but are less than 1% different after only 10 runs.

TOUR DESTINATION CHOICE

The San Francisco model includes separate tour destination choice models for each tour purpose as well as purpose-based intermediate stop choice models. In addition to considering two geographic scales (neighborhood and TAZ), this analysis considers destination choice for all purposes as well as workplace destination choice only. The analysis also considers the variation associated with tour making destined for a given neighborhood or TAZ as well as origin–destination pair variation, such as neighborhood to neighborhood.

Figure 5 presents a summary of tour destinations (or attractions) from all origins for all purposes to Neighborhood 10, an employment-rich area of the city, as indicated in Table 1. It demonstrates very consistent results, in that all the sample means are less than 0.2% different than the final mean tour destinations to the neighborhood.

Figure 6 presents a summary of tours between a set of neighborhood origins and destinations for all purposes. This represents a rarer choice in that only a single pair of origin and destination neighborhoods is considered—from Neighborhood 6 to Neighborhood 10. Accordingly, greater differences between the run and final means are observed. However, these differences are still quite small and demonstrate the stability of the destination choice models at this level of (dis)aggregation. The sample mean is never more than 1.5% different from the final mean.

Figure 7 presents a summary of tours between a set of neighborhood origins and destinations for work purposes only. Given that this is a rarer choice still, one would expect the differences observed to be greater than for all purposes and to require more runs to approach the final mean. Interestingly, although the work purpose differences appear slightly less consistent than the all-purpose difference as indicated by the greater standard error, the difference is more tightly bounded—never exceeding a 1% difference from the final result after 100 runs.

TOUR MODE CHOICE

The San Francisco model has two distinct mode choice models. The first model, the tour mode choice model, predicts the primary mode for a given tour and includes six choices. The second model, the trip mode choice model, predicts the mode for each individual trip segment of a tour and includes 11 choices. This paper presents the results of the tour mode choice models only but includes an analysis of mode choice by origin as well as by destination. For ease of presentation, the tour modes are aggregated into three classes: automobile, transit, and nonmotorized.

Figures 8 through 10 present the results of origin-based mode choices at the three geographic levels, meaning that mode for tours originating in the area are presented. Figure 8 presents the results of the origin-based mode choice analysis for automobile tours. These results demonstrate the continuing stability of the model at all levels of geographic aggregation. The difference in automobile tours at the county and neighborhood levels in the sample means is never more than 1% of the final mean. The TAZ level shows greater variation, but within 10 runs the difference at the TAZ level is less than 0.5%. Also, note the similarities between Figures 3 (percent zero-vehicle households), Figure 4 (trips per person), and Figure 8 (automobile tours).

Figure 9 presents the same analysis applied to transit tours. Again, the county- and neighborhood-level results are very stable; within five runs, the mean settles to a 1% difference from the final mean. Greater variation is observed at the TAZ level. This is likely a result of the slightly lower choice probabilities associated with transit for this particular origin zone—in absolute numbers there are about 4 times as many automobile tours as transit tours. While the TAZ-level differences in transit tours fluctuate more than TAZ-level differences in automobile tours, note that from the second run the difference is less than 2% from run 100.

Figure 10 presents this origin-based mode choice analysis applied to nonmotorized trips. These results are very similar to the transit results, which makes sense given that the choice probabilities for nonmotorized tours are very similar to those for transit, and the populations for which these probabilities are calculated are identical. However, it does appear that there is slightly greater variation in nonmotorized tours, as indicated by the greater spread of differences and the greater number of runs before the differences are less than 2%.

Figure 11 presents a slightly different view of the results. Instead of presenting the three geographic scales for a single mode, the origin-based summary of three modes is presented for a single geographic scale—a single TAZ. It indicates that all modes are within 2% of the final mean after only 10 runs; it also demonstrates a relative consistency across modes before the means appear to be consistent.

Figure 12 presents a destination-based analysis of mode choice, meaning that the modes for tours destined to the area are presented, and demonstrates that overall stability persists. Within 10 runs, the estimates for all modes are within 5% of the final mean. A comparison of Figure 12 (destination-based TAZ mode) and Figure 11 (origin-based TAZ mode) suggests that the origin-based mode choices overall are more stable than destination-based choices, as indicated by the smaller scale of percent differences from the final mean and by the faster convergence to the mean.

MULTIPLE SAMPLES

While the previous analyses indicate high levels of stability across all model components, would subsequent random samples converge to the same mean, and would this convergence occur at the same rate? To address these questions, the tour generation results from the 100 runs were divided into four sets of 25 runs each. The running means for trips per person were then calculated for each set and compared with the original final mean from 100 runs. Figure 13 presents the results of this analysis at the TAZ level, where one would expect to observe the greatest amount of variation and the greatest number of runs to converge. It was observed earlier that trips per person at the TAZ level converge to within 1% of the final 100-run mean after only 10 runs (Figure 4). Figure 13 indicates that this relatively rapid convergence persists across the four smaller samples—within 10 runs, all are within 1% of the final mean after 100 runs. Figure 13 also indicates that these smaller samples are converging to approximately the same final mean.

CONCLUSIONS

All the model components demonstrated high levels of stability, even when summarized at the smallest geographic levels or when interested in rarer choices. A relatively small number of runs are required at any geographic scale for confidence in the stability of the results. At the county and neighborhood levels, across all model components, the result of the initial run is often only a fraction of a percent different from the result after 100 runs. An implication of stability at the county level and neighborhood level of aggregation is that it can provide model developers with confidence that proper calibration of microsimulation models at aggregate levels does not require multiple runs, given that most model calibration would not focus on TAZ-level statistics. In addition, this aggregate stability can provide practitioners with confidence that overall measures generated by these models are reliable.

However, there is danger in running models only once and expecting that the indications yielded by that run are reasonably close to a global mean or expected value at anything other than the most aggregate levels. Table 1 presents descriptive statistics related to the figures in this paper. The table shows that, for the 100 runs that were conducted as part of this analysis, the minimum or maximum value at the TAZ level for an individual run could be as much as 10% to 25% different from the final mean. It is possible that some other random number sequences could yield larger percentage differences.

It is highly relevant in practice to understand how many runs are necessary to ensure confidence at different geographic resolutions. In general, the more specific the policy or statistic of interest, the smaller the geographic level, or the rarer a particular choice of interest, the greater the number of runs required. However, it is difficult to identify explicit criteria for the number of runs required to achieve a proper convergence or to even establish explicit convergence criteria. Any threshold would depend on the project analysis needs, the policies being tested, the particular statistics of interest, and the comfort level of the ultimate decision makers. As a result, the authors avoid proposing acceptable levels of convergence or the number of runs required to meet convergence. However, the plots show a trend, consistent with probability theory and the central limit theorem, that the number of iterations required is inversely proportional to the selection probability and the number of agents making a particular selection.

The authors also note that the results presented are particular to the San Francisco tour-based models. Although the general trends and indications are presumably applicable to other tour-based microsimulation models, specific instances may yield different results. It would be useful to conduct analyses similar to those presented here with other model systems, both to examine the transferability of the conclusions and to provide analysis specific to those models for future reference as they are used in application.

ACKNOWLEDGMENTS

The authors acknowledge the many contributors to the design, development, and implementation of the San Francisco model as well as the San Francisco County Transportation Authority staff who conceived the idea of developing an activity-based model for San Francisco. First, the authors thank the consultant project manager for model development, Maren Outwater of Cambridge Systematics, who guided the project through its completion. Also thanked is Jose Luis Moscovich of the San Francisco County Transportation Authority, who developed the initial concept and succeeded in funding the initial efforts. The consultant team includes those involved in developing the data (Kevin Tierney, Kazi Ullah, Erin Vaca, Rob Hranac, Steve Decker, Frank Goetzke, and Louisa Yue of Cambridge Systematics; Sally Nielson of Hausrath Economics Group; Mike Aronson of CCS; and Ross Maxwell of Parsons Brinckerhoff), those involved in developing the model (Nageswar Jonnalagadda, Earl Ruiter, and Arun Kuppam of Cambridge Systematics; William Davidson and Jim Hicks of Parsons Brinckerhoff; T. Keith Lawton of Lawton Consulting; and John Douglas Hunt of University of Calgary), and finally those involved in implementing the model (Vamsee Modugula of Cambridge Systematics and Larry Seiders of Urban Analysis Group).

REFERENCES

1. Veldhuisen, J., H. Timmermans, and L. Kapoen. Microsimulation Model of Activity-Travel Patterns and Traffic Flows: Specification, Validation Tests, and Monte Carlo Error. In *Transportation Research Record: Journal of the Transportation Research Board, No. 1706*, TRB, National Research Council, Washington, D.C., 2000, pp. 126–135.
2. Bowman, J. L., and M. E. Ben-Akiva. Day Activity Schedule Approach to Travel Demand Analysis. Presented at 78th Annual Meeting of the Transportation Research Board, Washington, D.C., 1999.
3. Bradley, M., M. L. Outwater, N. Jonnalagadda, and E. R. Ruiter. Estimation of Activity-Based Microsimulation Model for San Francisco. Presented at 80th Annual Meeting of the Transportation Research Board, Washington, D.C., 2001.
4. Jonnalagadda, N., J. Freedman, W. A. Davidson, and J. D. Hunt. Development of Microsimulation Activity-Based Model for San Francisco: Destination and Mode Choice Models. In *Transportation Research Record: Journal of the Transportation Research Board, No. 1777*, TRB, National Research Council, Washington, D.C., 2001, pp. 25–35.

Publication of this paper sponsored by Committee on Passenger Travel Demand Forecasting.