APPENDIX 11B – FREEWAY FACILITY CALIBRATION

The HCM 6th Edition emphasizes the importance of calibrating the freeway facility method to local traffic conditions, and reflecting locally observed free-flow speeds, bottleneck capacities, and congestion patterns. Calibration is most applicable at the facility level, because segment methods are limited to the analysis of undersaturated (i.e., under-capacity) conditions. The HCM provides guidance on recommended calibration steps, which have been adapted here for application in Oregon.

The calibration guidance in this appendix is intended to be used for the detailed analysis (HCM core methodology) method and for reliability analysis. For broad brush or screening methods, typically no detailed calibration is performed. However, the analyst should conduct a reasonableness check even for those methods and evaluate the method results against expected performance and study hypotheses. Even without formal calibration steps, the discussion in this appendix can be useful to inform the reasonableness checks for broad brush and screening methods.

Detailed Analysis Method Calibration

The recommended Oregon-specific freeway facility calibration method for detailed analyses consists of five steps as outlined below, and discussed in the following sections.

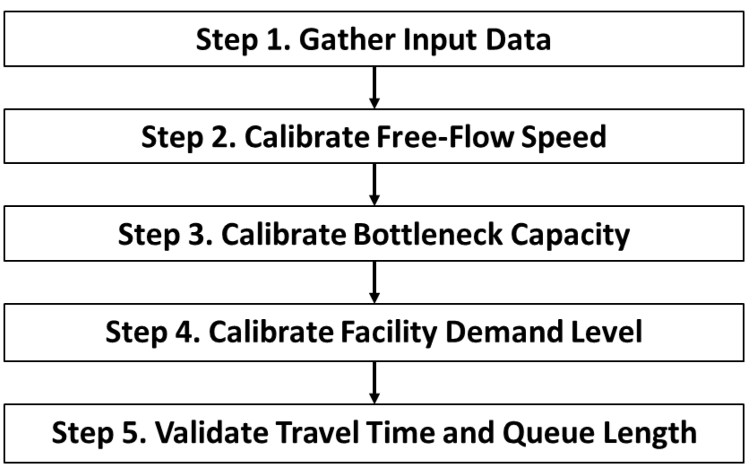

Step 1: Gather Input Data

Gather all input data required for a single freeway facility analysis. These data include:

1. Geometric information, including segment type, segment length, number of lanes, and relevant ramp information.

2. Facility free-flow speed (FFS)—see Appendix 11A for measurement and estimation techniques.
3. Estimated bottleneck capacity, from field data or local default values. In particular, capacity may need to be adjusted in merge, diverge, and weaving segments. Chapter 26, Section 5 in the HCM 6th Edition provides a detailed procedure for estimating bottleneck capacity from sensor data.
4. Demand-level data for all segments in all time intervals of the study period.
5. Facility performance data, including 15-minute travel times and queue lengths.

Step 2: Calibrate Free-Flow Speed

FFS can be field-measured and input directly, or estimated as described in Appendix 11A. Further calibration of the FFS can be applied in either case; however, if accurate field measurements of FFS are available, care should be taken before changing a field-measured input.

The analyst should select a time interval with a demand level sufficiently below capacity and with no active bottleneck, so that “free-flow” conditions can be captured. In a later step, the analyst will look at congested time periods; therefore, the study period should be chosen such that it includes free-flow conditions both before and after the congested time periods. This process also allows for “warm-up” and “cooldown” periods that help to ensure the build-up and dissipation of congestion are modeled accurately.

The analyst should also consider the driver population using the facility. If a significant portion of the drivers using the facility during the study period are unfamiliar with the route, capacity and the FFS may be lower than in the situation where nearly all drivers are familiar with the facility. The CAFs and SAFs for different driver populations recommended by the HCM are shown in Exhibit 1, along with Oregon-specific examples of situations where these factors might be applicable.

Exhibit 1 Capacity and Speed Adjustment Factors for Driver Population

Level of Driver Familiarity	CAF_{pop}	SAF_{pop}	Examples
All familiar drivers, regular commuters	1.000	1.000	I-105, I-205, OR 217
Mostly familiar drivers	0.968	0.975	I-5 (Willamette, Rogue Valleys)
Balanced mix of familiar and unfamiliar drivers	0.939	0.950	Many rural highways
Mostly unfamiliar drivers	0.898	0.913	OR 46 (Oregon Caves)
All or overwhelmingly unfamiliar drivers	0.852	0.863	Solar eclipse, emergency detour

Source: Derived from HCM 6th Edition, Exhibit 26-9.

The calibration process proceeds with the analyst using a computational engine to perform the analysis for the calibration day. The calibration day is a representative day with recurring congestion on the facility, without any atypical weather, work zone, or incident events. For a standard weekday analysis, the calibration day should be either a Tuesday, Wednesday, or Thursday.

To calibrate the free-flow speed, the analyst determines the predicted average travel time for a low-demand time interval and compares it to the observed travel time. The analyst then needs to repeatedly perform one of the following actions until the predicted facility travel time is within a predefined threshold (e.g. 10% error tolerance) of the observed facility travel time:

- *Reduce* the FFS in 1- to 5-mi/hr increments if the predicted travel time is *less* than the observed travel time, or
- *Increase* the FFS in 1- to 5-mi/hr increments if the predicted travel time is *more* than the observed travel time.

This process should only be used for analysis periods with demand levels far less than oversaturation (i.e., free-flow conditions).

It is important to note that a global FFS may not be appropriate for all facilities. FFS can vary between segments for several reasons, including varying speed limits or constrictive geometric conditions (e.g., narrow lanes, presence of a tunnel). To account for these conditions, the analyst can manually input FFS to match observed values or alternatively can use additional SAFs to adjust the segment's base FFS for the entire study period or even for specific analysis periods.

Step 3: Calibrate Bottleneck Capacity

In this step, the location and extent of bottlenecks are calibrated. This step requires a freeway facility to feature at least some periods of oversaturated flow conditions. Guidance for selecting capacity measurement locations and for reducing the collected data is provided in Chapter 26 of the HCM. Local default capacity values (see Appendix 11C) should be applied before embarking on significant data collection at bottlenecks.

It is very important to calibrate a facility's capacities, as the controlling capacity at the bottleneck is often significantly less than the HCM's base capacity. Three parameters are used to calibration for the location and extent of bottlenecks:

1. *Pre-breakdown capacity* at the bottleneck, implemented through a CAF relative to the freeway segment's base capacity. The HCM defines the pre-breakdown flow rate as the 15-minute average flow rate immediately prior to the breakdown event. For the purposes of this appendix, the pre-breakdown flow rate is equivalent to the segment capacity.
2. *Queue discharge rate* at the bottleneck following breakdown, implemented through a percentage capacity drop α. The HCM defines the queue discharge rate as the average flow rate during oversaturated conditions (i.e., during the time interval after breakdown and prior to recovery).
3. *Jam density* of the queue forming upstream of the bottleneck, which describes the maximum density (minimum inter-vehicle spacing) in a queued condition.

The pre-breakdown capacity and the queue-discharge capacity loss influence the actual throughput of the bottleneck, as well as the speed of shock waves describing the rate of change of the back of the queue. Jam density does not affect throughput; it only influences the formation and dissipation of queues at a bottleneck.

To calibrate for bottlenecks, the analyst needs to change the capacity and capacity drop values for different segments of the freeway facility to recreate the bottlenecks that are observed in the field. This is done using a CAF.

If these initial values predict the bottleneck location correctly, the analysis proceeds to the validation step. If the model fails to identify a bottleneck, the analyst should reduce capacity in increments of 50 pc/h/ln until a bottleneck occurs. However, if the HCM model identifies a bottleneck that does not exist in the field, the analyst should increase capacity in increments of 50 pc/h/ln until the bottleneck disappears.

It is recommended that the analyst wait to adjust the capacity drop value until after the bottleneck locations have been fixed. This procedure is performed as part of validating the queue length and travel time, as explained in Step 5.

Step 4: Calibrate Facility Demand Level

The demand level is a model input that can serve as a calibration parameter of last resort. Demands presumably have been determined from field data, and therefore can be considered as a fixed input. However, given the variability of demand (i.e., day-to-day fluctuation), as well as potential errors in volume and demand measurements (e.g. demand metering from upstream bottlenecks), demand can become a calibration parameter after the FFS and capacity adjustment possibilities have been exhausted.

Two potential problems may be encountered with demand levels. First, in oversaturated conditions, it is not possible to measure the demand level downstream of a bottleneck or within a queued segment. The *volume served* is measured, rather than the true segment demand. Second, demand data vary from day to day, and the selected demand levels may not represent a "typical" day. This second problem is also true if AADT demand values are used to estimate peak period demands. As a result, although demand level is one of the inputs to the core freeway facility analysis, it may be subject to calibration.

The analyst should increase the demand level in increments of 50 pc/h/ln until all bottlenecks that are observed in the field are activated in the freeway facility core analysis. However, if the model predicts bottlenecks that do not exist in the field, the user should decrease the demand level in increments of 50 pc/h/ln until those bottlenecks are deactivated. This activity should be performed in conjunction with Step 3: Calibrate Bottleneck Capacity.

Step 5: Validate Travel Time and Queue Length

The validation step has two major components:

1. Validate facility travel time, and
2. Validate queue length at active bottlenecks.

Travel Time Validation

After fixing the FFS and the bottleneck locations, the analyst should adjust the calibration parameters further to match predicted and observed facility travel times within a defined range (a 10% or less difference is recommended) during non-free-flow conditions. Note that FFS has already been fixed in Step 3 and will not be adjusted further in this step. This process can be done by adjusting

1. Demand level,
2. Pre-breakdown capacity,
3. Capacity drop, and
4. Jam density.

Probe datasets providing speeds and/or travel times for individual segments can be used to inform the validation. Any probe data set should be processed to generate 15-minute profiles of facility travel time targets, as well as a complete set of target segment speed contours.

If the model underestimates the travel time, the analyst should consider one of the following actions:

1. Increase the demand level (in increments of 100 pc/h/ln),
2. Reduce pre-breakdown capacity (in increments of 100 pc/h/ln), or
3. Increase the capacity drop (in increments of 1%).

If the model overestimates travel time, the analyst should consider one of the following actions:

1. Reduce the demand level (in increments of 50 pc/h/ln),
2. Increase pre-breakdown capacity (in increments of 50 pc/h/ln), or
3. Reduce the capacity drop (in increments of 1%).

Note that jam density is unlikely to have a significant impact on facility travel time and is therefore not included in the steps above.

Queue Length Validation

After the facility travel time is fixed, the queue lengths at the facility's active bottlenecks are expected to generally match field-observed conditions (i.e., within 10%). However, if differences exist, adjustments can be made to the capacity drop and/or jam density.

If the predicted queue length at an active bottleneck is *shorter* than observed in the field, the capacity drop should be *increased*, and the jam density should be *decreased*.

However, if the predicted queue length is *longer* than that observed in the field, the capacity drop should be *decreased*, and the jam density should be *increased*. It is recommended that the capacity drop be changed in increments of 1% and that the jam density be changed in increments of 10 pc/mi/ln.

When specific field observations of queue length are unavailable, this process can be informed by inspecting probe speed and travel time data at the TMC level (or sub-TMC level, if available, to more accurately identify the extent of queueing). Although aggregated facility travel times are useful for travel time validation, utilizing data at a level closer to the segment level is more appropriate for queue length validation. As congestion spills back from a bottleneck or breakdown, the individual TMC speeds and travel times will reflect the queue formation and dissipation, providing estimates of both spatial and temporal bounds. A visual comparison of speed contours and the associated extent of congestion can inform the queue length validation.

It is important to note that there is no guarantee of consistency between HCM segment boundaries and TMC boundaries. In fact, the HCM segmentation guidance will often result in segments smaller than those provided by most probe data vendors. Therefore, the HCM's predictive method may provide a higher-resolution look at queue length, resulting in more "realistic" queue formation and dissipation patterns than the averaged behavior of the probe data, which can appear overly blocky.

Freeway Facility Example Calibration

To illustrate the calibration of a freeway facility, the following analysis is presented for a facility with 34 segments that were analyzed (and calibrated) over an entire 24-hour period. While many freeway facility analyses are likely to focus on one peak period (e.g. AM peak or PM peak) this example illustrates the full capabilities of the method. A schematic of the facility is shown below in Exhibit 2.

Exhibit 2 Calibration Example Facility Geometry

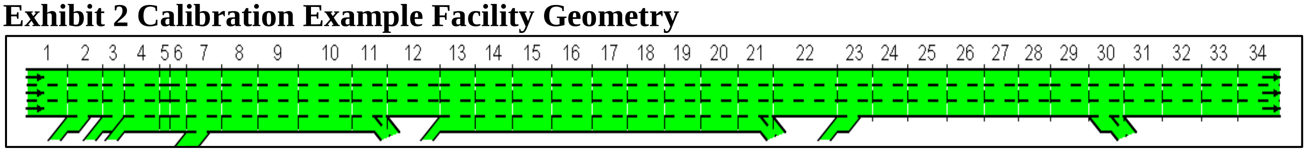

For the sample facility, travel times and speed data were collected from a probe-based data source for a typical calibration day (only recurring congestion). Alternatively, multiple days of probe data can be collected, and results averaged. The empirical speed contours are shown on the left in Exhibit 3, while the calibrated freeway facility results are shown on the right. For each case, segments are shown in columns from left to right, while time periods are shown in rows from midnight to midnight for a 24-hour analysis. Green cells represent fast speeds, while red cells represent slow or congested speeds. The yellow speeds at the beginning of the facility on the left, are sections with slower free-flow speeds (a tunnel section), while segments in the middle and at the end of the facility have higher free-flow speeds.

Exhibit 3 Calibration Example Speed Contours

Field data **FREEVAL results**

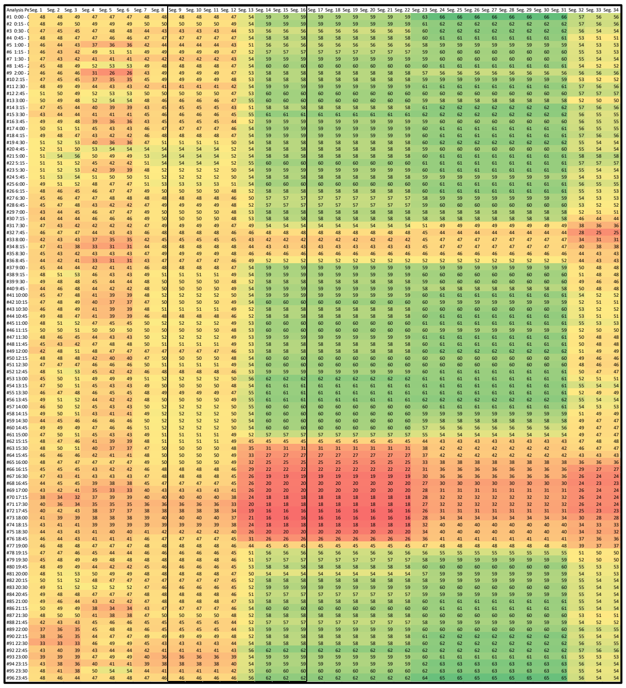

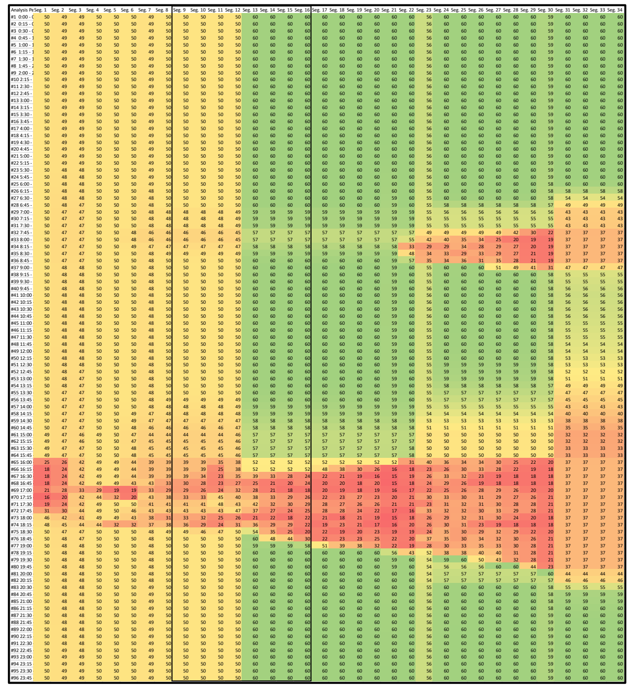

To evaluate the calibration results, 15-minute travel times can be compared between the field data, and the freeway facility results, as shown in Exhibit 4.

Exhibit 4 Calibration Example Travel Time Calibration Results (24 hours)

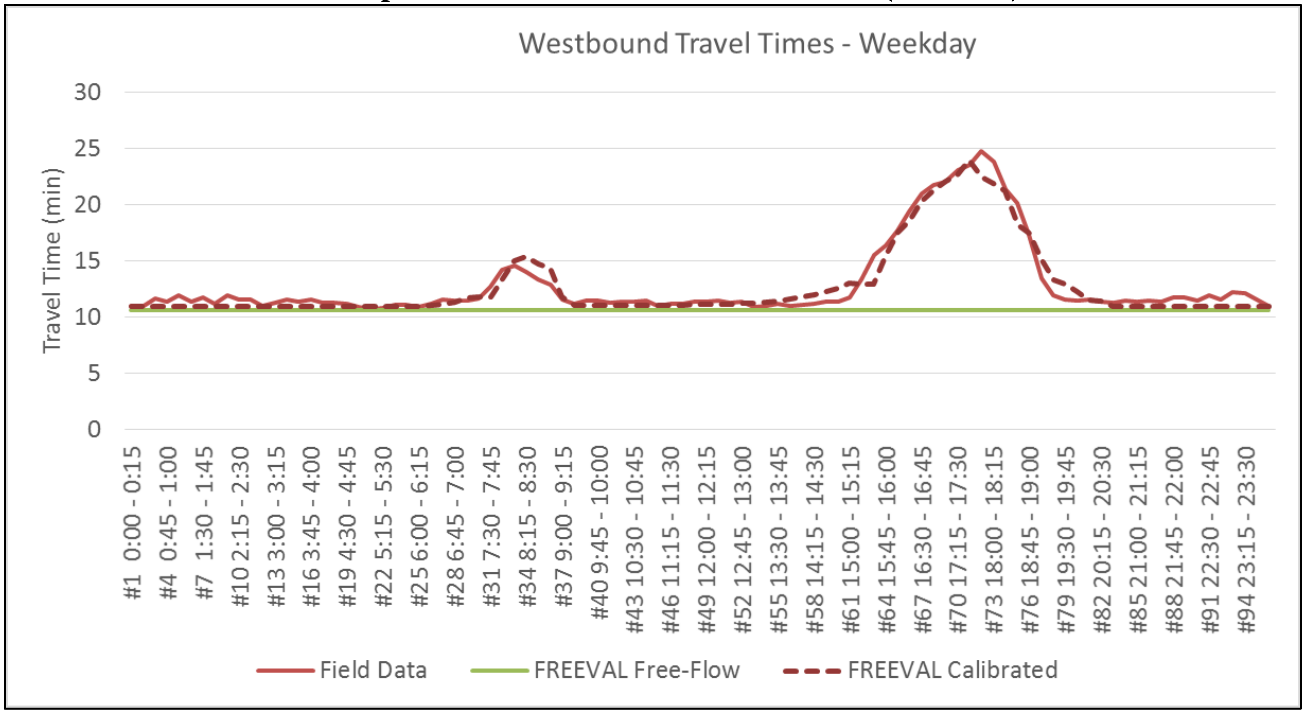

Facility Analysis and Additional Scenario Calibration

Once steps 1–5 have been completed, the analyst has completed the calibration of the facility to existing conditions. This base model can now be use to examine the impact of specific conditions or scenarios. For example, an analyst may wish to model a set of work zone scenarios to gauge how the facility performs under differing configurations such as a single-lane versus a two-lane closure.

While in many ways related to reliability analysis, this type of scenario analysis differs from the guidance of the reliability method in that it does not seek to model facility performance over the course of a year, but rather focuses on its performance under specific conditions. Some types of scenarios that can be investigated are: future demand forecasts, work zones, crash events, and facility performance under severe weather conditions. An analysis scenario can consist of one or any combination of these events, as long as the underlying assumptions required for the analysis remain the same as those of the calibrated base facility.

After deciding on the scenario parameters, the analyst should collect all relevant data related to changes in facility demand, capacity, speed, and geometry for the specific scenario. For example, when modeling a work zone scenario, the analyst should determine how the work zone configuration affects the existing segmentation (e.g., accounting for lane closures), whether a work zone speed limit will be in effect, any expected re-routing or traffic diversion that might occur, and any additional expected impacts on capacity (e.g., narrowed lanes or shifting traffic flows).

When conducting a scenario analysis, local calibration values should always be applied where available. However, the HCM provides some guidance and default values to assist in modeling certain types of scenarios. The reliability method presented in Chapter 11 in the HCM 6th Edition provides default CAF and SAF values for weather events based on free-flow speed (see Exhibit 5 and Exhibit 6), as well as default distributions of incident severity type and incident capacity adjustment factors (Exhibit 7 and Exhibit 8). Chapter 10 in the HCM 6th Edition also presents two models for capacity and free-flow speed adjustments specific to work zones.

Exhibit 5 Capacity Adjustment Factors by Weather Condition and Free-Flow Speed

Weather Type	Weather Event Definition	Free-Flow Speed (mph)				
		55	60	65	70	75
Medium rain	>0.10 – 0.25 in./h	0.94	0.93	0.92	0.91	0.90
Heavy rain	>0.25 in./h	0.89	0.88	0.86	0.84	0.82
Light snow	>0.00 – 0.05 in./h	0.97	0.96	0.96	0.95	0.95
Light-medium snow	>0.05 – 0.10 in./h	0.95	0.94	0.92	0.90	0.88
Medium-heavy snow	>0.10 – 0.50 in./h	0.93	0.91	0.90	0.88	0.87
Heavy snow	>0.50 in./h	0.80	0.78	0.76	0.74	0.72
Severe Cold	<−4° F	0.93	0.92	0.92	0.91	0.90
Low visibility	0.50 – 0.99 mi	0.90	0.90	0.90	0.90	0.90
Very low visibility	0.25 – 0.49 mi	0.88	0.88	0.88	0.88	0.88
Minimal visibility	<0.25 mi	0.90	0.90	0.90	0.90	0.90
Non-severe weather	All conditions not listed above	1.00	1.00	100	1.00	1.00

Source: HCM 6th Edition, Exhibit 11-20.

Exhibit 6 Speed Adjustment Factors by Weather Condition and Free-Flow Speed

Weather Type	Weather Event Definition	Free-Flow Speed (mph) 55	60	65	70	75
Medium rain	>0.10-0.25 in./h	0.96	0.95	0.94	0.93	0.93
Heavy rain	>0.25 in./h	0.94	0.93	0.93	0.92	0.91
Light snow	>0.00-0.05 in./h	0.94	0.92	0.89	0.87	0.84
Light-medium snow	>0.05-0.10 in./h	0.92	0.90	0.88	0.86	0.83
Medium-heavy snow	>0.10-0.0 5 in./h	0.90	0.88	0.86	0.84	0.82
Heavy snow	>0.50 in./h	0.88	0.86	0.85	0.83	0.81
Severe Cold	<-4° F	0.95	0.95	0.94	0.93	0.92
Low visibility	0.50 – 0.99 mi	0.96	0.95	0.94	0.94	0.93
Very low visibility	0.25 – 0.49 mi	0.95	0.94	0.93	0.92	0.91
Minimal visibility	<0.25 mi	0.95	0.94	0.93	0.92	0.91
Non-severe weather	All conditions not listed above	1.00	1.00	100	1.00	100

Source: HCM 6th Edition, Exhibit 11-21.

Exhibit 7 Default Distributions and Durations of Different Incident Types

Parameter	Incident Severity Type: Shoulder Closure	1 Lane Closed	2 Lanes Closed	3 Lanes Closed	4+ Lanes Closed
Distribution (%)	75.4	19.6	3.1	1.9	0
Duration, minutes (mean)	34	34.6	53.6	67.9	67.9
Duration (std. dev.)	15.1	13.8	13.9	21.9	21.9
Duration (min)	8.7	16	30.5	36	36
Duration (max)	58	58.2	66.9	93.3	93.3

Source: HCM 6th Edition, Exhibit 11-22.

Exhibit 8 CAFs by Incident Type and Number of Directional Lanes on the Facility

Directional Lanes	Incident Severity Type					
	No Incident	Shoulder Closure	1 Lane Closed	2 Lanes Closed	3 Lanes Closed	4+ Lanes Closed
2	1.00	0.81	0.70	N/A	N/A	N/A
3	1.00	0.83	0.74	0.51	N/A	N/A
4	1.00	0.85	0.77	0.5	0.52	N/A
5	1.00	0.87	0.81	0.67	0.50	0.50
6	1.00	0.89	0.85	0.75	0.52	0.52
7	1.00	0.91	0.88	0.80	0.63	0.63
8	1.00	0.93	0.89	0.84	0.66	0.66

Source: HCM 6th Edition, Exhibit 11-23.
Note: The methodology does not permit all lanes to be closed.

Estimating 15-Minute Volumes from Hourly Volumes

In order to create 15-minute volumes in the event that (1) only hourly volumes are available and (2) the study period is longer than the period of time for which traffic volumes are available, the following guidance is applicable:

1. When only hourly volumes are available, the 15-minute hourly flow rates can be determined through linear interpolation. The following simple example demonstrates how this works for two consecutive hours where the hourly demand values are 1000 vph and 2000 vph, respectively:

    a. Period #1 (6:00 am) = 1000 + 0 × (2000−1000)/4 = 1000
    b. Period #2 (6:15 am) = 1000 + 1 × (2000−1000)/4 = 1250
    c. Period #3 (6:30 am) = 1000 + 2 × (2000−1000)/4 = 1500
    d. Period #4 (6:45am) = 1000 + 3 × (2000-1000)/4 = 1750
    e. Period #5 (7:00 am) = 1000 + 4 × (2000−1000)/4 = 2000

2. When the study period is longer than the period of time for which traffic volumes are available, the user can use a profile that best represents the facility to determine additional demand values. The values can be estimated based on that profile using either an AADT value or matching known values to percentages of the profile in the respective periods. The same linear interpolation procedure described in the previous bullet can then be used to estimate 15-minute flow rates.

Reliability Method Calibration

This section provides guidance on calibrating the reliability analysis method for freeway facilities as defined in the HCM 6th Edition. As with the core freeway facility method, it is important to calibrate the analysis to local conditions. To the extent possible, a reliability analysis should be consistent with known existing or historical conditions before being used for predictive analysis. Reliability calibration should only be started *after* completing the calibration at the core facility level, as described earlier in this appendix.

The overarching goal of calibration for the reliability method is to produce a representative distribution of the travel times experienced by motorists. However, instead of comparing actual travel times, it is recommended that travel time indexes (TTIs) be compared during the calibration process. TTIs are calculated as the predicted travel time divided by the free-flow travel time, and therefore represent a more normalized, unitless performance metric for the facility. While accuracy thresholds can vary for specific applications depending on data availability and other considerations, the general calibration target is for the 50^{th}, 80^{th}, and 95^{th} percentile TTI values of the estimated distribution to match within 10% of field observations.

Calibration of higher percentile travel times (e.g. 98^{th} or 99^{th} percentile) is not recommended, since these events are due to very rare outlier events. Conceptually, the 80^{th} percentile corresponds to the travel times expected once a week (1 in 5 days), while the 95^{th} percentile is roughly equivalent to the travel times experiences once a month (1 in 20 days).

The calibration guidance in this appendix is largely intended for the detailed reliability method. For screening- (planning-) level methods, no detailed calibration is generally necessary. However, even for simplified methods, it is important for the analyst to at least verify the reasonableness of the results against expected performance. Even if a formal calibration process is not conducted, the guidance provided in this section can help in performing reasonableness checks for screening-level methods.

The reliability calibration process involves five steps (numbered 0–4), as illustrated below:

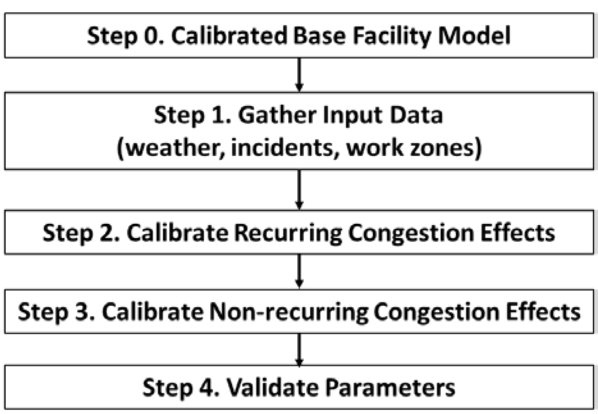

Step 0: Obtain Calibrated Base Facility File

This initial item is a prerequisite for conducting a reliability analysis using the HCM methodology. For guidance on the creation of a core facility analysis, refer to APM Section 11.3.4. Guidance on calibrating a core model was provided earlier in this appendix.

As previously emphasized, it is important to first calibrate the core facility model for most applications. However, while a well-calibrated core facility model provides the best starting point for an HCM-based reliability analysis, there is no guarantee that it will automatically produce a well-calibrated facility reliability analysis. Therefore, the following steps and validation processes should be conducted for any detailed analysis.

Step 1: Gather Reliability Input Data for the Reporting Period

The first step of conducting a reliability analysis is to set the desired *reliability reporting period* (RRP). The RRP consists of the days of the year that will be used to measure or forecast the facility's reliability. APM Section 9.3.4 illustrates the RRP. The days included in the reliability reporting period will depend on the type of facility being analyzed and the analysis purpose. For example, a study of the reliability of a major commute route within a metropolitan area might define a reliability reporting period of all non-holiday weekdays during the year. In contrast, a study of the reliability of a highway leading from the Willamette Valley to the Oregon Coast might define a reliability reporting period consisting of Saturdays, Sundays, and holidays during the summer. Other analyses might use all the days in a year.

With an RRP specified for the project, the analyst should gather the necessary input data required for a reliability analysis over the time period. These data are described in APM Section 11.5.4, Step 2b. Next, the analyst can perform an initial scenario generation run to estimate the TTI distribution for the facility. If the resulting distribution matches the field-observed distribution within acceptable thresholds, the analyst need not pursue additional calibration. However, if the estimated distribution does not match the field-observed distribution to an acceptable level, the analyst should proceed with Steps 2 through 4 until the calibration targets are achieved.

Step 2: Calibrate Recurring Congestion Effects

When travel times greater than the free-flow travel time occur on a facility, the underlying sources fall into one of two groups: (1) recurring congestion due to normal heavy volumes, or (2) non-recurring congestion due to events such as incidents or severe weather.

Consider the field-observed and estimated TTI distributions in Exhibit 9. The estimated distribution does not align with the field-observed distribution. When discrepancies occur in the body of the distribution (e.g., from the 1^{st} through the 50^{th} - 60^{th} percentile), it likely indicates that the effects of recurring congestion on the facility are not being modeled accurately. In cases where discrepancies occur in the tail of the distribution, the issue likely lies with non-recurring congestion effects. In this example, discrepancies occur in both the body and the tail.

Exhibit 9 Example Comparison of Field-Observed and HCM-Estimated TTI Distributions

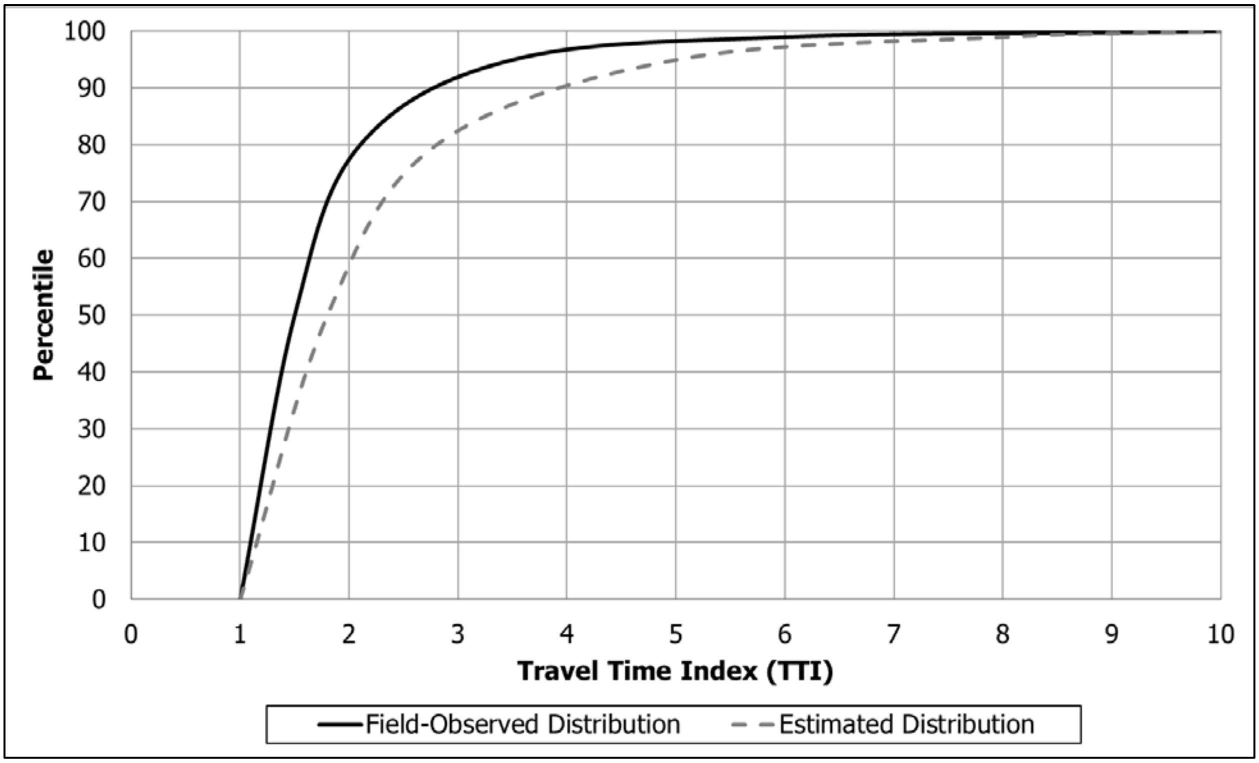

Source: HCM 6th Edition, Exhibit 25-33.

The demand level for the seed day used by the core facility method is set before the start of a reliability analysis, as it is determined from initial input data and adjusted as needed during the core facility calibration process. However, a facility's demand almost always experiences significant seasonal variability, requiring the use of monthly and day-of-week demand multipliers. While the calibration of the seed day to observed data may be accurate, it is possible there was an unusual or unaccounted-for variation in demand that conflicts with one or more of the reliability demand multiplier inputs.

For the distribution shown above in Exhibit 9, the estimated TTI distribution does not match the field-observed distribution for almost any percentile. Furthermore, the estimated distribution is shifted, showing estimated travel times are almost always worse than actual conditions. This result indicates that the demand level in the seed file may be too high in the context of the reliability seasonal demand multipliers. To correct this situation, the analyst can lower the demand level of the seed file before conducting additional reliability analysis runs.

Alternatively, if the estimated distribution is largely inaccurate in the body and shifted to the left of the field-observed data, this is an indication that the demand level on the seed day may be too low. An example of this type of behavior is shown in Exhibit 10. To correct this issue, the analyst can raise the demand level of the seed file before performing additional reliability runs.

Exhibit 10 Example Comparison of Field-Observed and Estimated TTI Distributions where Recurring Congestion Effects Are Likely Underestimated

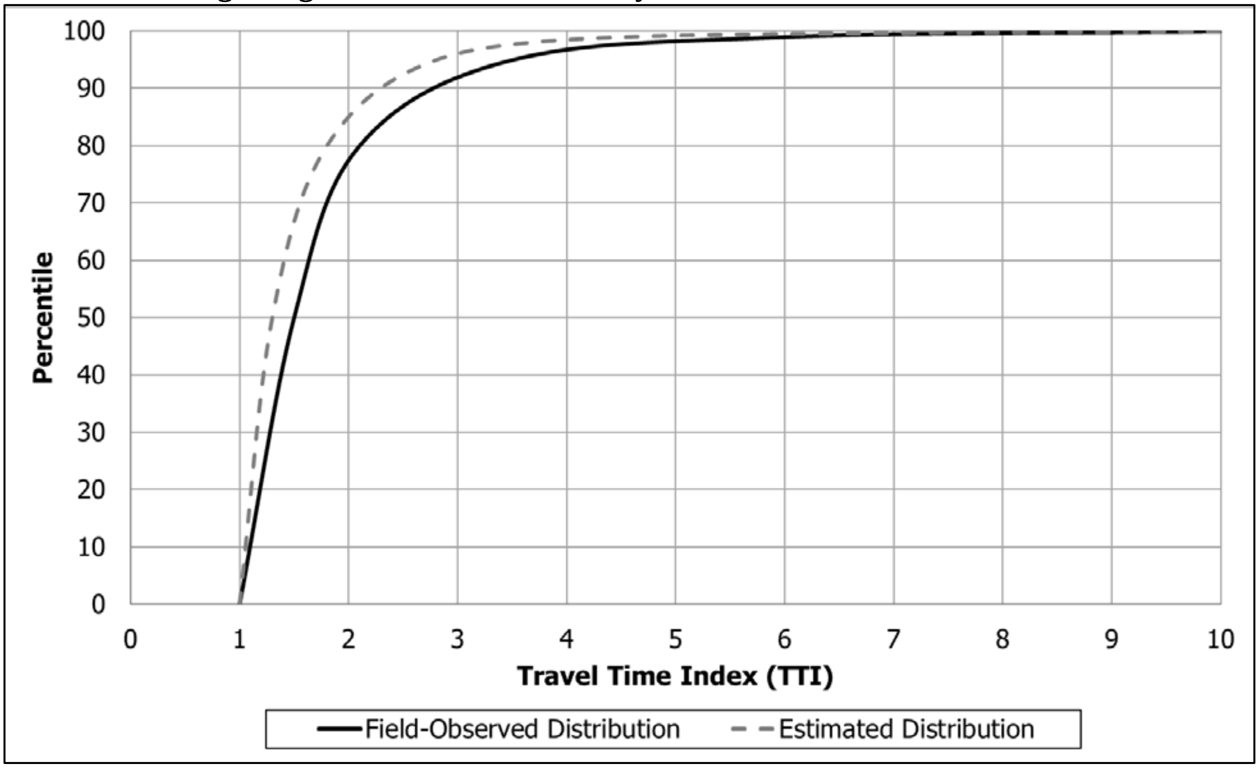

Source: HCM 6th Edition, Exhibit 25-34.

In both Exhibit 9 and Exhibit 10, the *x*-axis intercept is very similar for the estimated and field-observed distributions. Since the estimated and field data are normalized using the same free-flow travel time, this result indicates that the travel times experienced at very low volumes are being accurately modeled. However, if the two distributions are not similar at very low flow rates, it is likely that there is an issue with the initial free-flow speed calibration step for the core method.

Step 3: Calibrate Non-recurring Congestion Effects

Once the demand level has been calibrated, the distributions of estimated and field-observed TTI should match acceptably well at least through the 50th - 60th percentile. As the higher-percentile TTI values are generally caused by non-recurring congestion, it is possible that adjusting the demand level will not fully address the differences. Exhibit 11 provides an example of what the comparison of distributions might look like if non-recurring congestion events are causing travel times to be overestimated.

Exhibit 11 Example Comparison of Field-Observed and Estimated TTI Distributions where Differences Are Largely Observed in the Higher Percentiles

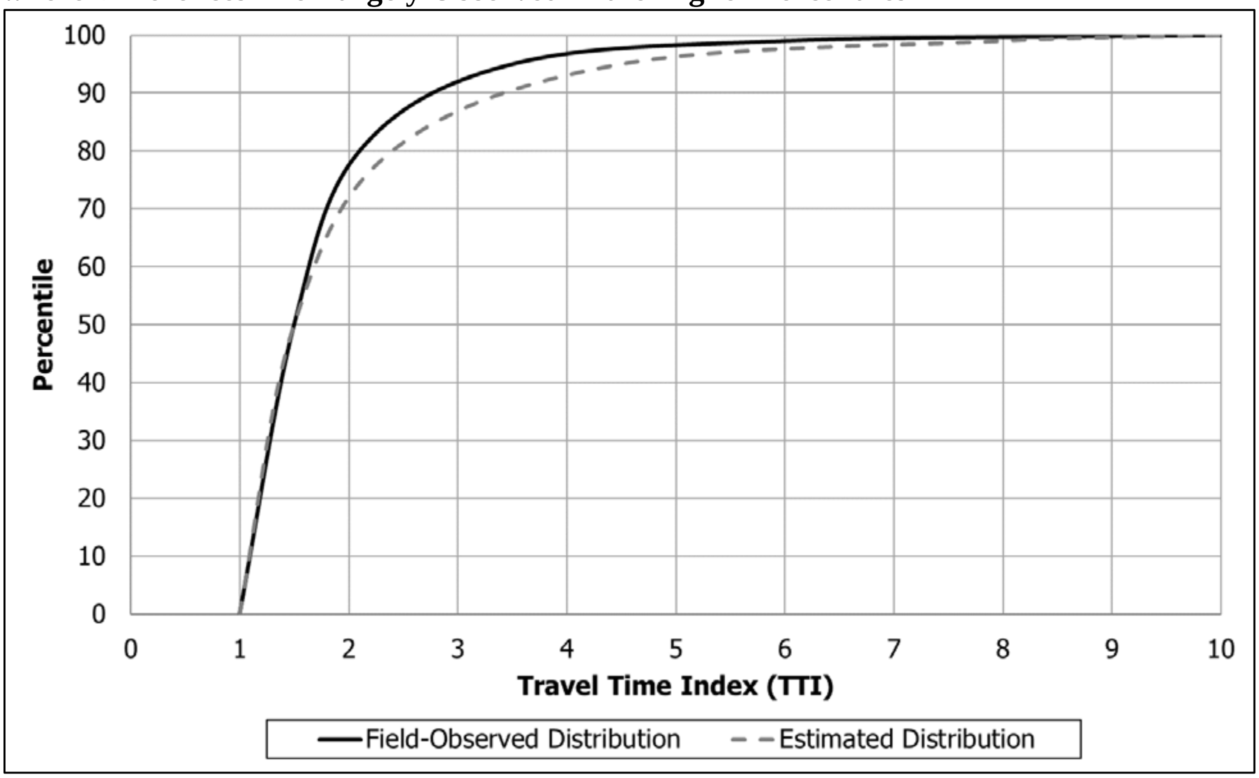

Source: HCM 6th Edition, Exhibit 25-35.

Non-recurring congestion events fall into two primary categories: incidents and weather events. In almost all cases, incidents have a larger impact on reliability and congestion than weather events, and therefore should be the initial focus when calibrating a reliability analysis. Furthermore, weather data are likely to be more representative of actual averages than crash data, because historical weather data are more widely available. For example, it is recommended that weather event probabilities be developed using ten years of observational data, whereas gathering similarly extensive crash data is more challenging. At the same time, it is possible that 10 years of weather data could capture unusually extreme or rare weather conditions that may need to be screened out during the initial development of the weather input data.

Calibration parameters for adjusting the overall number and types of incidents consist of:

1. The monthly incident frequency per 100 million vehicle-miles,
2. The incident-to-crash ratio, and
3. The underlying probability distribution of incident severity.

Calibration parameters for adjusting the impacts of incident severities consist of:

1. Attributes determining the duration of incidents by severity type (mean duration, standard deviation, and duration range); and
2. Speed, capacity, and demand adjustments by severity type.

Reducing the frequency of crashes and incidents, or reducing the probability of severe incidents, should have the effect of reducing the number of travel times occurring in the tail of the estimated distribution. (Similarly, increasing either of these parameters would increase the number of travel times in the tail of the distribution). When adjusting the parameters of individual incident severities, reducing the demand adjustment or increasing the capacity or speed adjustments should also result in lower travel times (again, an increase in travel times can be achieved by making the opposite adjustment).

If calibrating incidents is not sufficient to produce a suitable match in the distributions, or if the analyst feels the incidents are accurately modeled, weather events can also be calibrated using the following parameters:

1. Monthly probability of each weather event type;
2. Weather event type durations; and
3. Speed, capacity, or demand adjustment factors associated with each event type.

Decreasing the probability or duration of weather events, reducing the demand adjustment, or increasing the capacity or speed adjustments will all tend to result in lower travel times (and vice versa).

Step 4: Validate Parameters

Due to the compounding effects of each aspect of reliability analysis, adjusting more than one factor at a time (i.e., demand, incidents, or weather) can cause unanticipated outcomes. Consequently, it is highly recommended that the analyst only adjust one factor at a time before conducting an additional reliability run and validating the new estimated distribution against field-observed data. As mentioned previously, a general threshold to target for a calibrated reliability analysis is to bring the 50th, 80th, and 95th percentile TTI values to within 10% of field observations.

When reflecting on the relative importance of accurately matching certain TTI percentiles, it can be informative to consider what that percentile represents in the context of the length of the RRP. For example, when considering a one-year RRP, the 50th percentile TTI is experienced by drivers approximately half the time, while the 99th percentile TTI would only occur a few times per year. Furthermore, it is likely that the extreme values at the tail of the TTI distribution are the product of extreme or compounding incidents, or unusually severe weather. When making investment decisions based on reliability, the relative rarity of the high TTI values found at the tail of the cumulative distribution may not warrant heavy consideration.

Exhibit 12 provides a mapping from TTI percentile to the associated approximate number of days during a one-year period that travelers on a given directional roadway section would experience that TTI.

Exhibit 12 TTI Percentiles Translated to Approximate Annual Frequencies of Occurrence

Percentile	Approximate Frequency of Occurrence by Direction
50th	Every other day
80th	Once every five days, or once per work week
95th	Once every twenty days, or once per month (weekdays only)
99th	Once every 100 days, or only two to three weekdays a year