CASTR

User Manual

Cost and Schedule Total Risk

Thousands of scenarios. One defensible estimate. Integrated probabilistic cost and schedule risk analysis for highway construction, built on NHI 134205 principles.

Version 2.5.0 (April 26, 2026)

Table of Contents

Chapter 1: Getting Started

1.1 Overview

CASTR (Probabilistic Risk-Based Estimating) is a desktop application designed for probabilistic cost and schedule risk analysis of highway construction projects. The application implements Monte Carlo simulation with Gaussian copula correlation modeling to produce rigorous, defensible risk-based cost and schedule estimates.

The methodology is based on the National Highway Institute (NHI) Course 134205, which provides a standardized framework for risk-based estimating in transportation projects. CASTR automates this methodology, making it accessible to project teams without requiring specialized statistical software.

Key capabilities of CASTR include probabilistic cost and schedule modeling, correlation analysis between risk variables, market conditions assessment, inflation adjustment to year-of-expenditure dollars, sensitivity analysis through tornado charts, and tail driver identification for understanding extreme outcomes.

1.2 Installation & Setup

System Requirements

- Operating System: Windows 10 or later (64-bit)
- Memory: 4 GB RAM minimum; 8 GB recommended for large simulations
- Disk Space: approximately 150 MB for the application
- Display: 1280×800 minimum resolution; 1920×1080 or higher recommended
- Browser: A modern web browser (Chrome, Edge, or Firefox) is used automatically for the user interface
- Network: No internet connection is required after installation. CASTR runs entirely offline on your local machine.

Installation Steps

CASTR is distributed as a self-contained executable and does not require a traditional installer. To install the application:

1. Download the CASTR application archive from your authorized distribution source.
2. Extract the downloaded archive to a folder of your choice (e.g., C:\Program Files\CASTR\ or C:\CASTR\).
3. Optionally, create a desktop shortcut to CASTR.exe for convenient access, or pin it to your Windows taskbar.
4. No additional software or database setup is required. CASTR creates a local SQLite database automatically on first launch to store your project data.

First Launch

When you run CASTR.exe for the first time, Windows Defender or your firewall may display a prompt asking you to allow local network access. Click Allow — this is required only for communication between the application and your web browser on your own machine. CASTR does not transmit any data over the internet.

Windows SmartScreen may also display a warning because the application is not yet code-signed. If you see this message, click “More info” and then “Run anyway” to proceed.

Updating CASTR

CASTR automatically checks for new versions each time the application is launched. If a newer version is available, an amber notification banner appears at the top of the screen displaying the new version number, a brief summary of changes, and a Download Update link. Clicking the link opens the download page in your web browser. The banner can be dismissed by clicking the close button; it will reappear the next time you launch the application until the update is installed.

To install an update, download the latest version and replace the existing CASTR.exe file in your installation folder. Your project data is stored separately in the local database and is not affected by updates. The update check requires an internet connection; if the application is offline, the check is silently skipped and the application continues to operate normally.

Uninstallation

To remove CASTR from your computer, simply delete the CASTR application folder. If you also wish to remove your project data, delete the local database file located in the same directory. CASTR does not modify the Windows registry or install system-level services, so no additional cleanup is needed.

Troubleshooting

- **Port 8000 already in use:** If CASTR fails to start, another application may be using port 8000. Close the conflicting application, or check Windows Task Manager for a previous CASTR instance that did not shut down properly.
- **Browser does not open automatically:** If your web browser does not open when you launch CASTR, manually navigate to http://localhost:8000 in Chrome, Edge, or Firefox.
- **Antivirus quarantine:** Some antivirus software may flag CASTR.exe because it is not code-signed. Add an exception for CASTR.exe in your antivirus settings if the application is blocked or quarantined.
- **Display scaling issues:** If the interface appears too small or too large, adjust your browser zoom level (Ctrl+Plus or Ctrl+Minus) to a comfortable size. The application is responsive and adapts to different zoom levels.

1.3 Launching the Application

To start CASTR, double-click the CASTR.exe file on your desktop or in the installation directory. The application will launch a local backend server and automatically open your default web browser to http://localhost:8000, where the user interface is displayed.

The backend server starts automatically and runs in the background. You do not need to install or configure any web server software. The application is entirely self-contained and all data is stored locally on your machine.

If the browser does not open automatically, you can manually navigate to http://localhost:8000 in any modern web browser (Chrome, Edge, or Firefox recommended).

1.4 License and Trial

CASTR includes a 90-day free evaluation period that allows you to assess the full functionality of the application. During the trial you are limited to a maximum of three projects, and a trial banner is displayed at the top of the application showing the number of days remaining. Output produced in trial mode is intended for evaluation only and should not be relied upon for commercial project estimates or bids.

After the trial period, you must purchase a license key to continue using CASTR without restrictions. Each license is priced at TBD and is bound to a single machine. To purchase, visit https://michaelsmith935-ux.github.io/castr-website/ and click Purchase License. During checkout you will be asked to provide the Machine ID shown in the top bar of your CASTR application. Your license key is delivered by email, typically within 24 hours of payment confirmation.

To activate your license, click the Trial chip in the top bar of CASTR and paste the key you received by email into the License Activation dialog. Once activated, the top bar will display a Licensed chip together with the date through which your installation is covered.

Upgrade Coverage Window. A new license includes twelve (12) months of free upgrades beginning on the date of purchase. During this period you may download and install any release of CASTR that is published within the window. When the coverage window expires, your existing installation and license key continue to work indefinitely without any further payment, but you will no longer be entitled to install releases published after the window expired.

Renewals. To regain access to newer releases after your coverage window expires, you may purchase a renewal at https://michaelsmith935-ux.github.io/castr-website/ for TBD. A renewal extends your Upgrade Coverage Window by an additional twelve (12) months and unlocks every release published during the renewed window. Renewal is optional: if you are satisfied with the version you are running, you never have to renew. Licenses purchased before 2026-04-14 include lifetime Upgrade Coverage at no additional cost and are not subject to the renewal model.

Coverage Banners. CASTR displays a banner at the top of the application when your coverage status changes. A blue banner appears when your coverage window is within 30 days of expiring, reminding you to renew if you wish to stay current. An amber banner appears after the window has expired if a newer release is available that you are not entitled to install. A red banner appears only in the rare case that you install a release published outside your coverage window; in that case, either renew your coverage or revert to a release that falls within your window.

Machine ID and Rehosts. License keys are tied to a specific Machine ID displayed in the top bar of the application. If you replace your computer, reinstall Windows, or upgrade hardware in a way that changes the Machine ID, you may request a rehost at no charge by emailing michael.smith935@gmail.com. Each license is entitled to up to two (2) rehosts per rolling 12-month period. When you request a rehost, you will be asked to confirm in writing that the original machine has been decommissioned or is no longer in use; the new key is bound to your new Machine ID and delivered by email, typically within one business day.

Refunds and Support. You may request a full refund of the license purchase price within thirty (30) days of purchase by emailing michael.smith935@gmail.com. Email support is included during your active Upgrade Coverage Window. The complete License Agreement is available at https://michaelsmith935-ux.github.io/castr-website/License.txt.

1.5 Application Layout

The CASTR interface is organized into three main areas: the sidebar navigation on the left, the top bar, and the main content area.

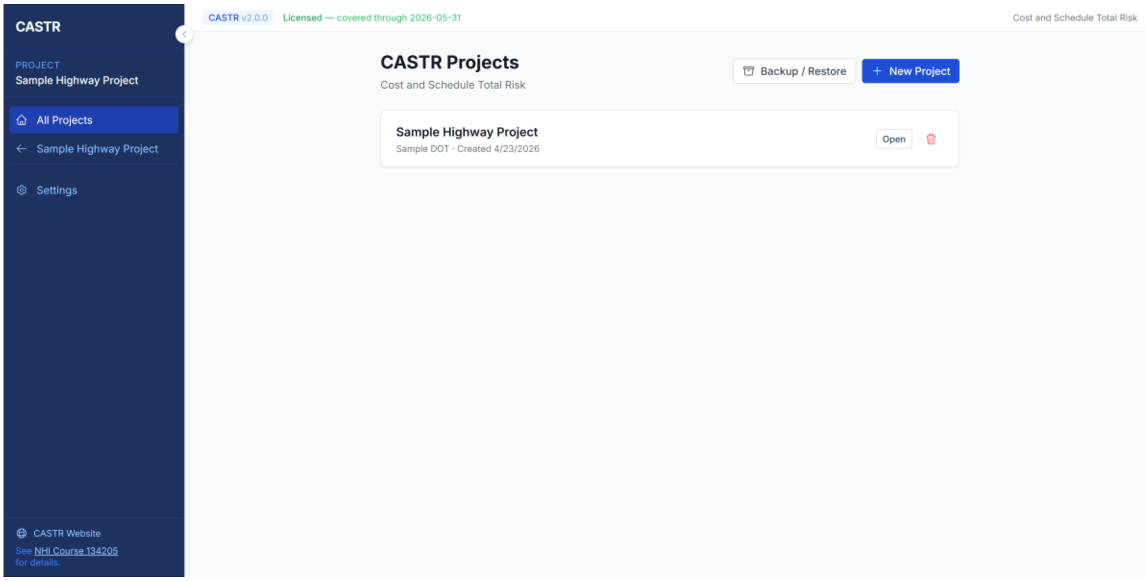

Figure 1.1 - Project List home page

The left sidebar navigation is collapsible and provides access to all pages within a project. The sidebar is organized into the following groups: Setup (Project Information, Segments and Phases, Inflation Tables); Inputs (Base Variability, Market Conditions); Risks (Risk Identification, Risk Register, Correlation Editor, Minor Risks, Review Staging); Analysis (Simulation, Results); Settings (Scoring Tables); and the Starter Library. When a project is open, a Progress Checklist appears in the sidebar showing completion status for key setup steps, helping guide new users through the workflow.

The top bar displays the application name, current version number, and trial status information. When a license key has been entered, the trial bar is replaced with a confirmation of your active license.

The main content area occupies the remainder of the screen and displays the currently selected page. All data entry and results viewing takes place in this area.

1.6 In-App Help System

CASTR includes a built-in help system with three complementary layers designed to assist both new and experienced users.

Tooltip Help Icons: Small question mark icons appear next to key field labels and column headers throughout the application. Click any icon to see a brief explanation of what that field means and how it is used. Tooltips close when you click elsewhere or press the Escape key.

Progress Checklist: When a project is open, the sidebar displays a Progress section that tracks completion of key setup steps such as setting the Estimate Last Update Date, defining segments and phases, creating inflation tables, adding base variability data, and running the simulation. Green checkmarks indicate completed steps, and clicking an incomplete step navigates directly to the relevant page. The checklist can be collapsed or hidden if desired.

Guide Panel: Every project page includes a Guide button in the page header. Clicking it opens a collapsible panel on the right side of the screen containing a tutorial for the current page. The panel explains the purpose of the page, describes key fields and actions, and provides tips for effective use. The content updates automatically as you navigate between pages. The panel can be closed at any time and its open/closed state is remembered across sessions.

FAQ Panel (added in v1.10): Every page also has an FAQ button next to the Guide button. Clicking it opens a collapsible Q&A panel on the right side. Each question expands to reveal a targeted answer. The FAQ is organized into themed groups (e.g., "Getting Started", "Troubleshooting") and is specific to the current page. Guide and FAQ panels are mutually exclusive - opening one closes the other - so the workspace stays uncluttered. Use the Guide for a top-to-bottom walkthrough of the page; use the FAQ when you have a specific "how do I..." or "why is this happening" question.

v2.2 Update — the in-app Guide and FAQ panels were rewritten to cover the new Inflation Sensitivity Scenarios workflow, the Probability Bands table on Scoring Tables, the Risk Impact tornado mode, the reconstructed-distribution view used by the SoSP / SoPSP budget percentile methods, the Word .docx export, and the Starter Library Excel import / Restore Defaults actions. Open Guide on any page after upgrading to read the per-page walk-through.

Chapter 2: Project Management

2.1 Creating a Project

To create a new project, navigate to the home page (Project List) and click the "+ New Project" button. A dialog will appear prompting you to enter the project name. After entering the name, the project is created and you are taken to the Project Information page where you can enter additional details.

Each project is stored independently in the application database. You can create as many projects as needed (subject to trial limitations for unlicensed installations).

2.2 Opening a Project

From the home page, each project is displayed as a card showing the project name, creation date, and a brief summary. Click the "Open" button on any project card to enter that project. Once a project is opened, the sidebar navigation expands to show all available project pages.

You can return to the home page at any time by clicking the CASTR logo or home icon at the top of the sidebar.

2.3 Duplicating a Project

To create a copy of an existing project, first open the project you want to duplicate. Navigate to the Project Information page and click the "Save As New Project..." button. This creates a complete copy of the project including all data, risk entries, correlations, inflation tables, and settings. The duplicated project will appear on the home page with the name you specify.

This feature is useful for creating scenario analyses where you want to compare different assumptions or risk profiles using the same base project data.

2.4 Deleting a Project

To delete a project, locate the project card on the home page and click the red trash icon. A confirmation dialog will appear to prevent accidental deletion. Once confirmed, the project and all associated data are permanently removed from the database.

> **Tip:** *Consider exporting a backup before deleting a project. Deleted projects cannot be recovered.*

2.5 Backup and Restore

CASTR provides a comprehensive backup and restore system accessible from the "Backup / Restore" button on the home page. This feature allows you to export projects for archival, transfer to another machine, or safeguard against data loss.

Export: On the Export tab, select one or more projects using the checkboxes, then click the download button. The selected projects are packaged into a single .zip file containing all project data.

Import: On the Import tab, upload a previously exported .zip file. The application will display a preview of the projects contained in the file. If any project names conflict with existing projects, you can choose a resolution strategy for each conflict: keep the existing project, replace it with the imported version, or rename the imported project. Once you have reviewed and resolved any conflicts, click Commit to complete the import.

Chapter 3: Project Setup

3.1 Project Information

The Project Information page is the first configuration step for any new project. It captures essential metadata about the project and the team members involved in the risk analysis.

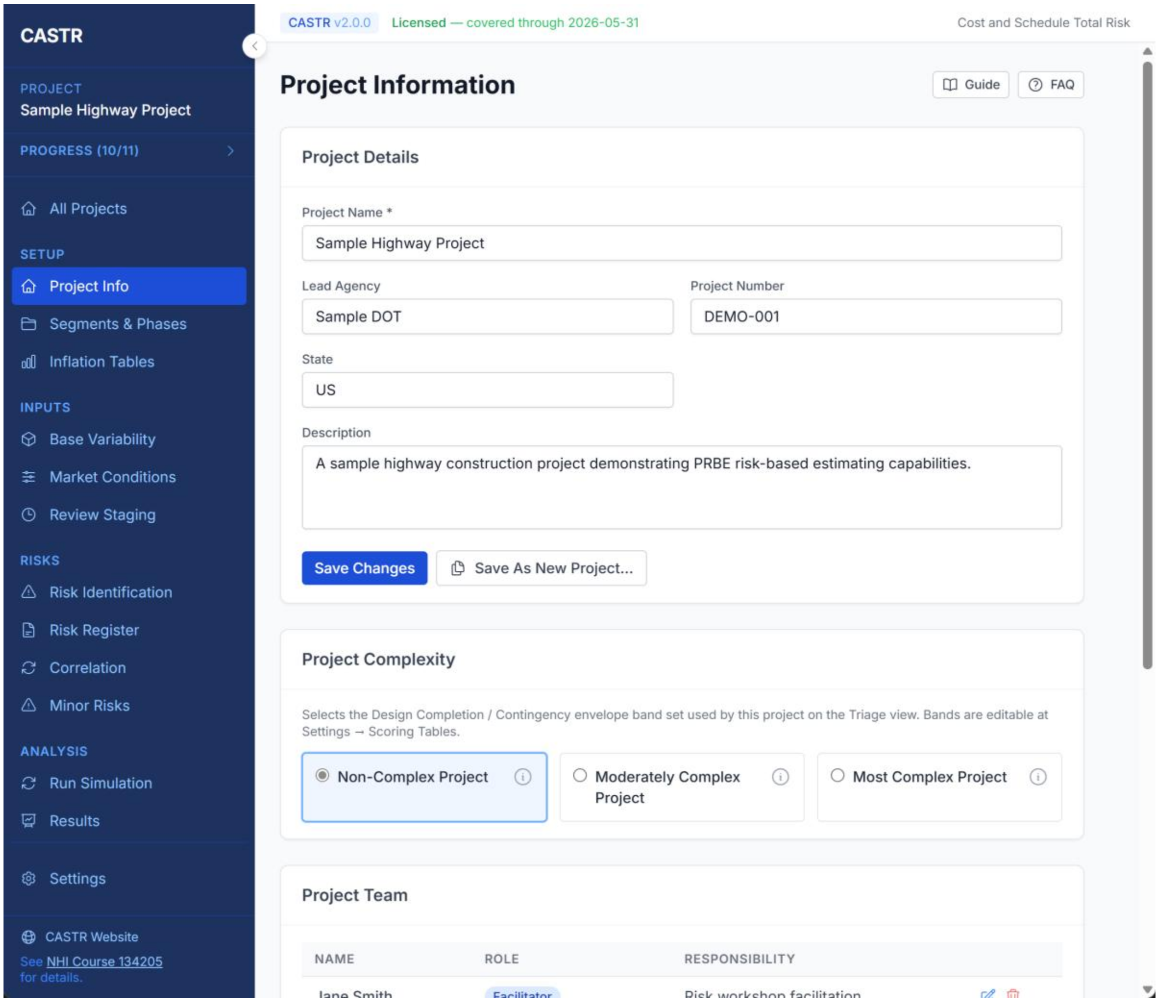

Figure 3.1 - Project Information page

Project Name is the only required field. It identifies the project throughout the application and in exported reports.

Additional fields include Lead Agency (the sponsoring organization), Project Number (your internal reference number), State (the U.S. state where the project is located), and Description (a free-text field for notes about the project scope or purpose).

Project Complexity: CASTR presents three complexity categories (Non-Complex, Moderately Complex, Most Complex) as an informational section below the Project Details card. The selection drives which Design-Completion contingency envelope band set is used by the Contingency Envelope panel on the Risk Identification page (Triage view). More complex projects legitimately warrant a wider reasonable contingency range at a given design stage. Click the info icon next to any complexity option to see a description of the typical characteristics of projects in that category; these definitions draw on the FHWA / NHI Course 134205 guidance. If the complexity you need is not a close match for any category, start with the middle band (Moderately Complex) and refine the envelope bands themselves at Settings → Scoring Tables. The selection may be changed at any time — only the Contingency Envelope reasonableness check is affected; existing risk data, costs, and distributions remain unchanged.

Team Management: Add team members who are involved in the risk analysis process. For each member, enter their Name, select a Role (Facilitator, Modeler, Observer, or Project Manager), and optionally describe their Responsibility. Team information is included in exported reports for documentation purposes.

Estimate Last Update Date: This is a critical field that serves as the baseline date for all inflation calculations. It represents the date when the project cost estimate was last updated. All inflation adjustments use this date as the starting point. Ensure this date is set correctly before running simulations that include inflation.

3.2 Segments and Phases

The Segments and Phases page defines the structural breakdown of your project. This structure is used throughout the application to organize cost estimates, risk entries, and simulation results.

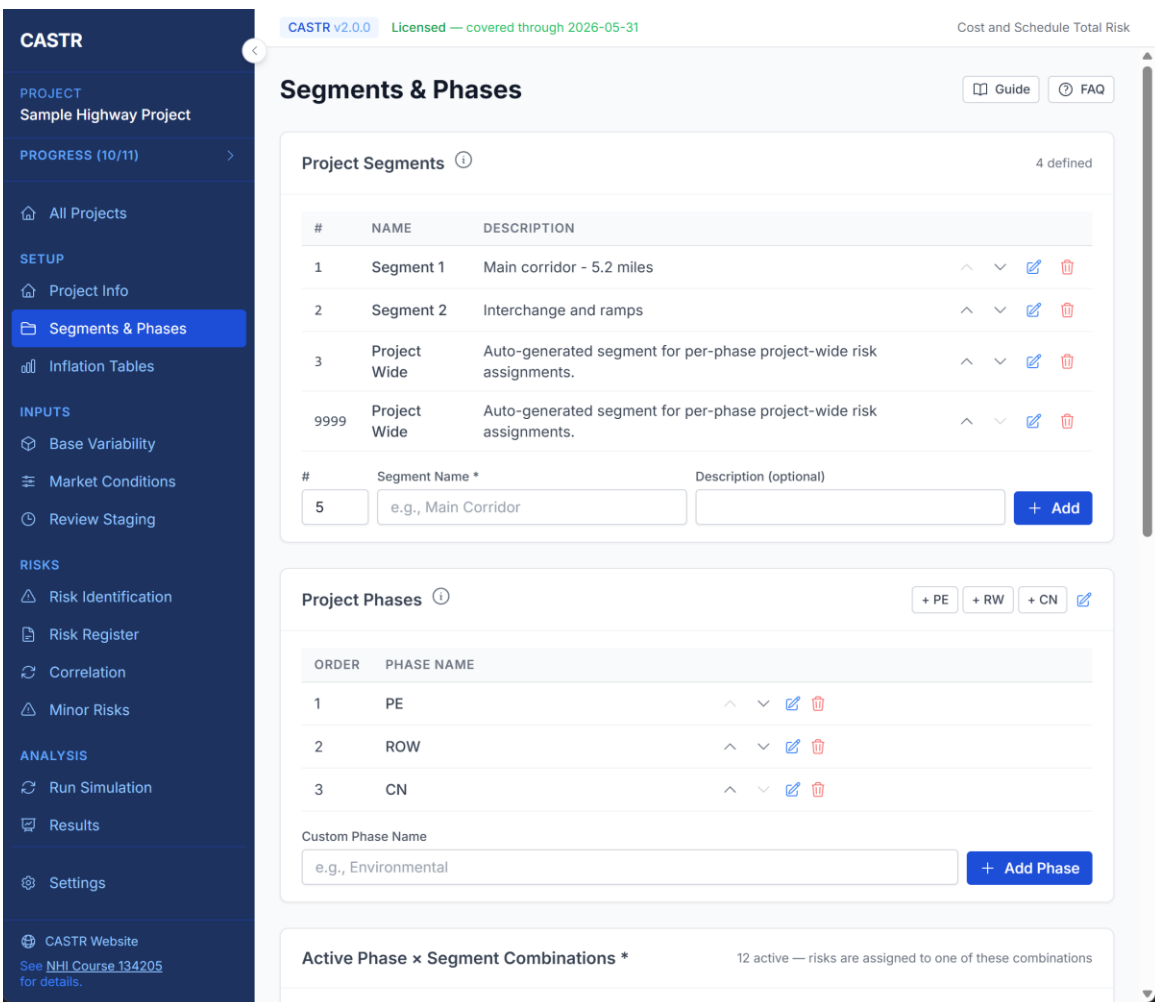

Figure 3.2 - Segments and Phases page

Segments divide the project geographically or logically. For example, a highway project might have segments such as "Main Corridor," "Interchange A," and "Bridge Structure." Each segment represents a distinct physical portion of the project.

Phases define the cost stages of the project. CASTR includes default phases: PE (Preliminary Engineering), ROW (Right-of-Way), and CN (Construction). Quick-add buttons are provided for these standard phases. You can also add custom phases as needed for your project.

Active Phase-Segment Combinations: A checkbox matrix is displayed showing all possible combinations of phases and segments. Check the boxes for the combinations that are applicable to your project. Only active combinations will appear in the data entry tables throughout the application.

Risk Categories: At the bottom of the page, you can define custom risk categories used to classify risks in the Risk Identification page. Common categories include ENV (Environmental), CON (Construction), DES (Design), ROW (Right-of-Way), and UTIL (Utilities). Categories help organize and filter risks during analysis.

3.3 Inflation Tables

Inflation tables define the annual inflation rates used to convert current-year costs to year-of-expenditure dollars. Each table contains a series of annual rates that are applied to costs based on their scheduled expenditure dates.

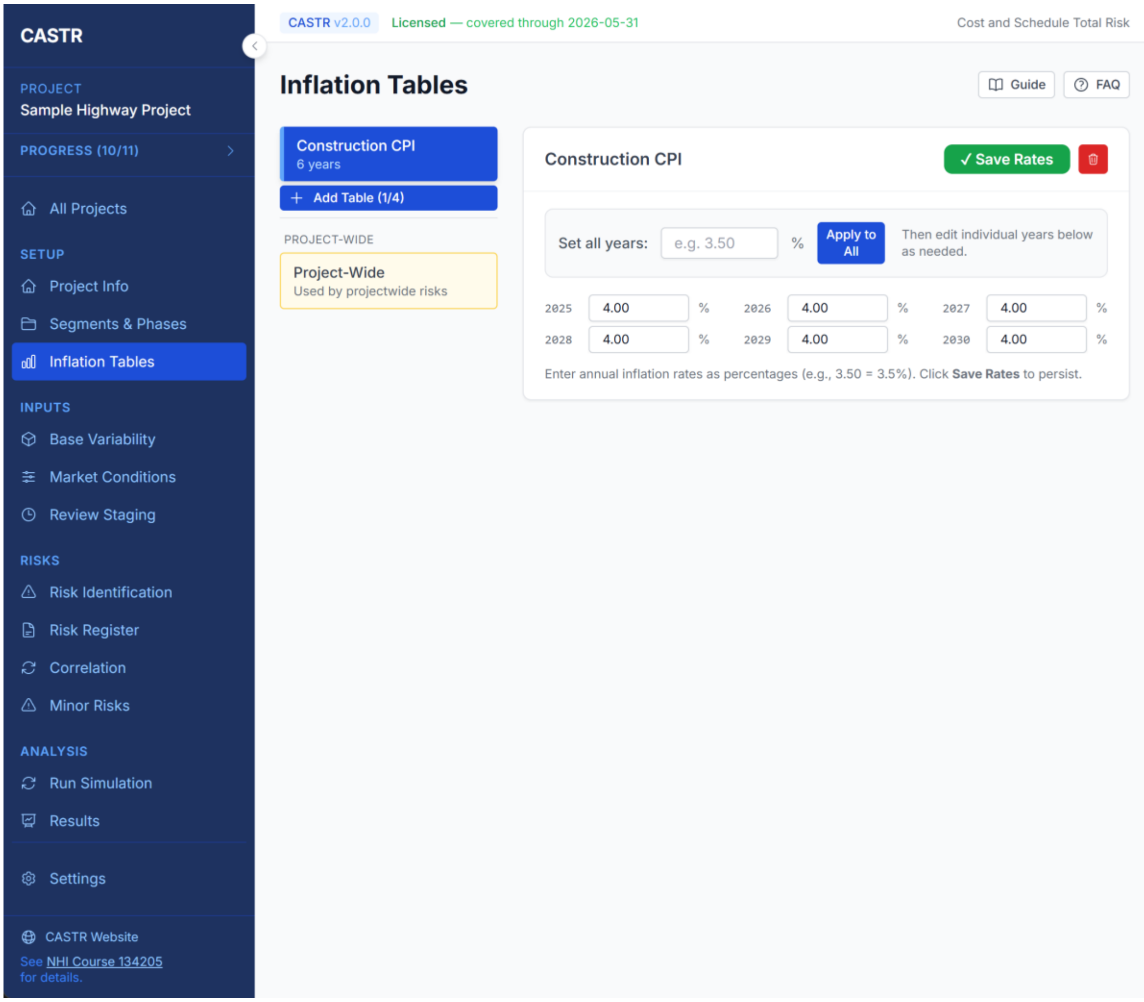

Figure 3.3 - Inflation Tables page

To create a new inflation table, click the "Add Table" button and provide a name for the table. Click on a table card to open it for editing, where you can enter year-by-year inflation rates. Each row represents a calendar year and its corresponding annual inflation rate as a percentage.

One inflation table can be designated as the Project-Wide table. This table is used for project-wide risks that are not assigned to a specific phase-segment. All other tables can be assigned to individual phases on the Segments and Phases page.

> **Tip:** *If your project uses a single inflation rate across all phases, you only need one table marked as Project-Wide.*

Chapter 4: Cost and Schedule Inputs

4.1 Base Variability

The Base Variability page is where you enter the foundational cost and schedule data for each phase-segment combination. This data forms the deterministic baseline upon which probabilistic analysis is applied.

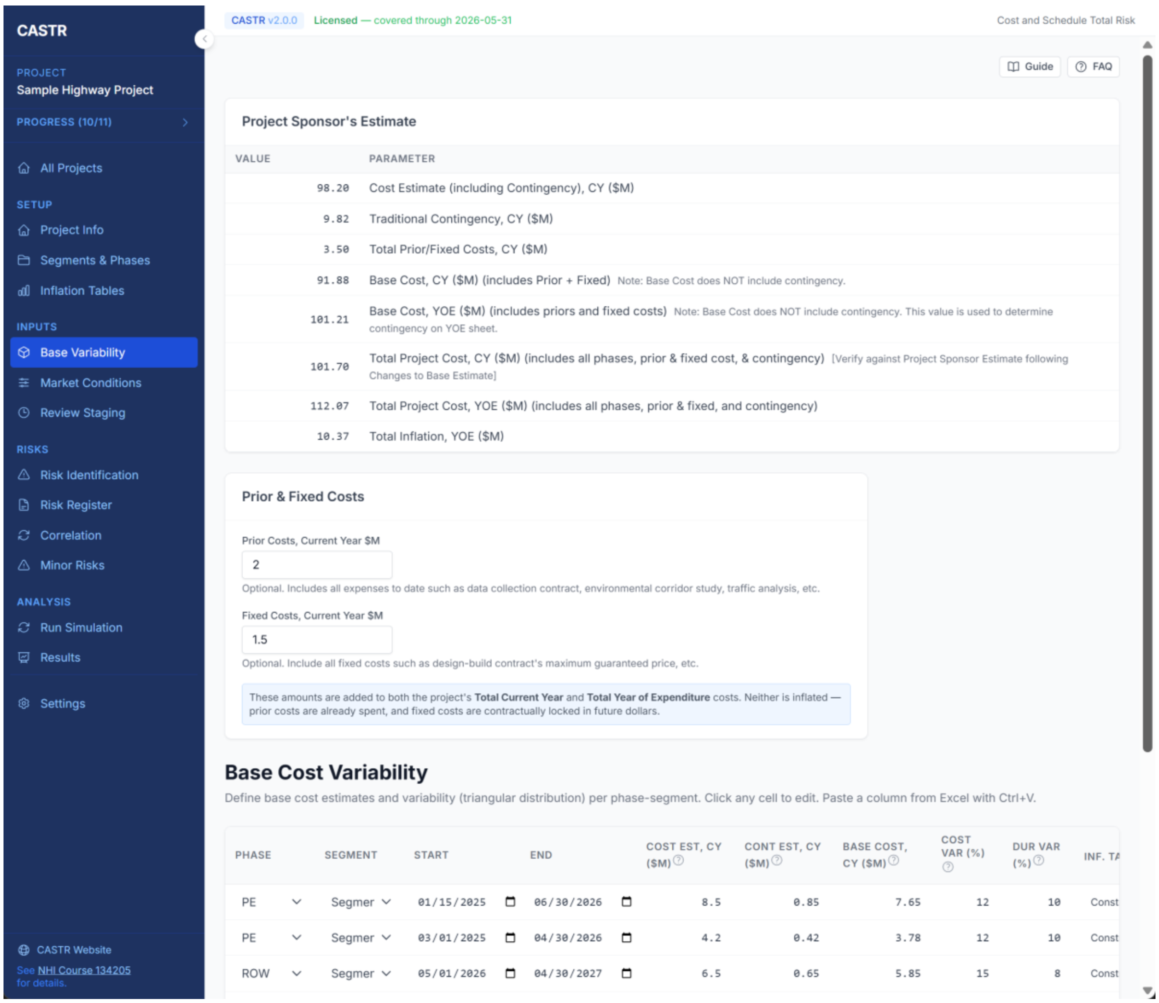

Figure 4.1 - Base Variability page

Project Sponsor's Estimate: The top section displays an auto-calculated summary of the total project cost based on the values entered in the table below. This provides a quick reference to the overall estimate as you enter data.

Prior and Fixed Costs: Prior costs represent sunk costs that have already been spent. Fixed costs are locked-in future expenditures that are not subject to variability. Neither Prior nor Fixed costs are inflated during year-of-expenditure calculations. These values are added to the total project cost but are not subject to probabilistic sampling.

Base Cost Variability Table: The main data entry table contains one row per active phase-segment combination. For each row, you enter the Start Date, End Date, Cost Est, CY (the current-year estimated phase-segment cost, including contingency for historical project risks —

do not include prior costs or fixed costs; enter those amounts above), Contingency Estimate (the current-year estimate of risks, often based on historical data), Cost Variability percentage, Duration Variability percentage, and the assigned Inflation Table. Base Cost CY is calculated automatically as Cost Est minus Cont Est and does not include contingency. Base Cost YOE is Base Cost CY multiplied by the inflation factor. Base Cost Infl is the amount of inflation computed from the Last Estimate Update date to the mid-point of the phase-segment. Cost Est YOE is the year of expenditure estimate including contingency.

Cost Variability and Duration Variability represent the inherent variability in the estimated base price and schedule of project elements. This uncertainty is neither influenced by nor caused by significant risk events that may or may not occur. For example, a cost variability of 10% means the simulated cost for that phase-segment can range from 90% to 110% of the base estimate, forming a symmetric triangular distribution. Usually, these values are small (i.e., ≤ 10%).

> **Tip:** *You can paste a column of values directly from Excel into the table using Ctrl+V. Click on the first cell in the target column before pasting.*

4.2 Market Conditions

The Market Conditions page models macro-level economic forces that affect all project costs. These conditions represent broad market trends such as labor market tightness, material price fluctuations, and competitive bidding environments. As of v1.9, market conditions are typically not applied to Preliminary Engineering (PE) phases because PE costs are staff and consultant hours rather than market-bid construction; the Sample Highway Project now reflects this by omitting PE phase-segments from its market-condition rows. New rows added via the “Add Row” button start completely blank: Phase and Segment appear as amber-highlighted “— Select —” dropdowns that the user must pick, and all multiplier and probability cells are empty. The triangular validator for each BTP / As-Planned / WTP scenario only fires once all three of Low, Most Likely, and High have been entered, so you can fill values in any order. Incomplete rows — those missing a Phase, a Segment, a multiplier, or any leg of a triangular triple — are skipped by the simulation engine without aborting the run.

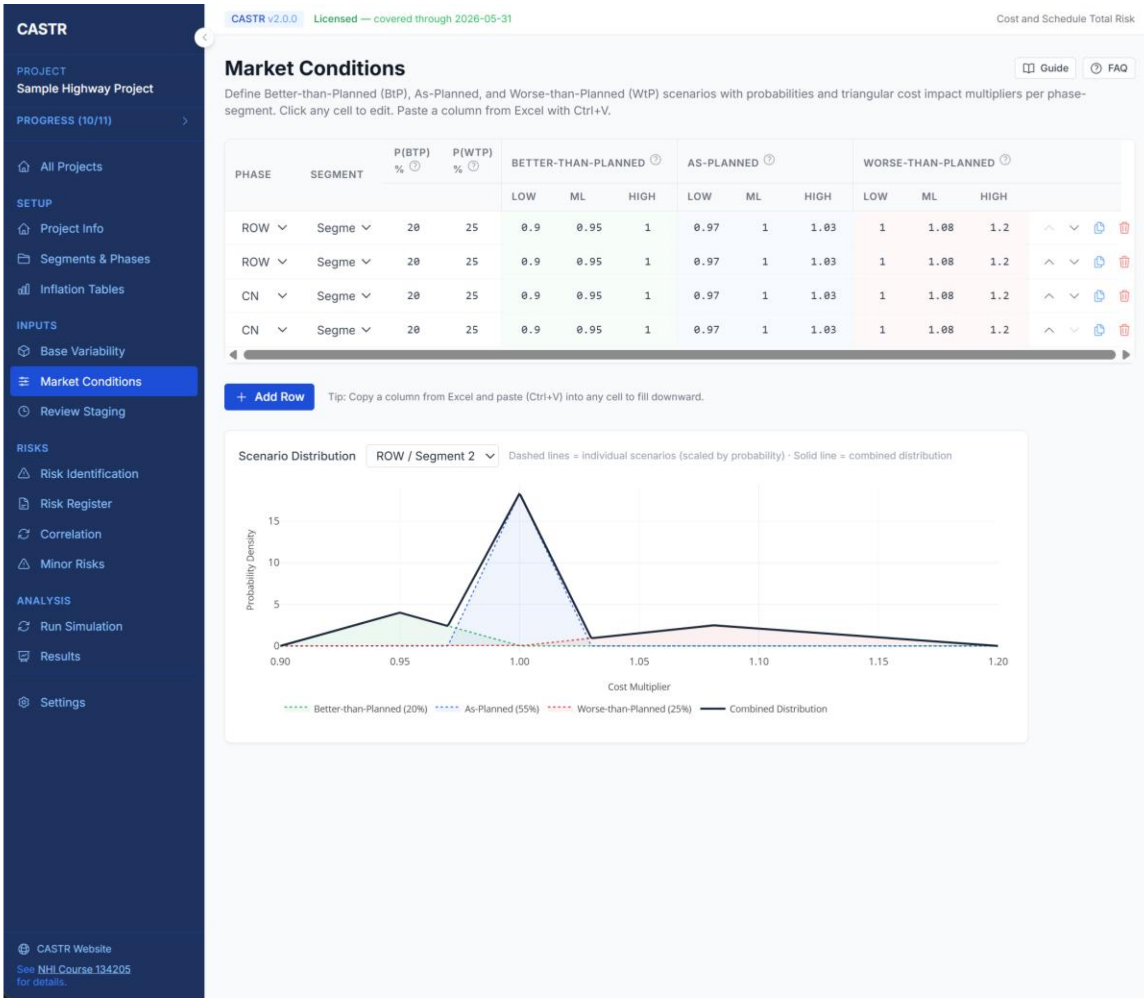

Figure 4.2 - Market Conditions page

For each phase-segment combination, three market scenarios are defined:

- **Better-Than-Planned (BTP):** The probability that market conditions at the time of project advertising and bid award will be more favorable than planned, resulting in bids lower than historical norms. Set the probability of this scenario occurring and define Low, Most Likely, and High cost multipliers (all less than 1.0).
- **As-Planned (AP):** The expected market condition. The probability is automatically calculated as 1 minus the BTP and WTP probabilities. Cost multipliers are centered around 1.0.
- **Worse-Than-Planned (WTP):** The probability that market conditions at the time of project advertising and bid award will be less favorable than planned, resulting in bids higher than historical norms. Set the probability and define cost multipliers (greater than 1.0).

Each scenario uses a triangular distribution defined by the Low, Most Likely, and High cost multipliers. During simulation, one scenario is selected based on the assigned probabilities, and the corresponding multiplier distribution is sampled and applied to all costs in that phase-segment.

4.2.1 Scenario Distribution Graph (added in v1.12)

Below the table is a Scenario Distribution chart that visualizes the mixture distribution for the phase-segment you select in the dropdown next to the chart title. It is a sanity-check of the values entered in the table row above — no new data is entered here. The Scenario Distribution chart visible at the bottom of Figure 4.2 (Market Conditions page) shows the CN / Seg 1 row.

How to read it: Three dashed curves show each scenario's triangular PDF scaled by that scenario's probability (area under each dashed curve equals that scenario's probability). The solid dark line is the combined mixture distribution — the sum of the three scaled curves — and represents what the simulation effectively draws from for this phase-segment. The x-axis is the cost multiplier (1.00 = no adjustment); the y-axis is probability density. Legend chips at the bottom show each scenario's name with its probability in parentheses.

How to use it: (1) Check where the solid line's bulk sits — if it peaks well above or below 1.00, you are effectively biasing cost up or down on that phase-segment. (2) Look for gaps between the dashed curves — gaps mean a range of multipliers has zero probability in any scenario (see §4.2.2). (3) Confirm the dashed peak heights reflect your probability allocation.

4.2.2 Gap Detection (added in v1.12)

CASTR flags distribution gaps in the Market Conditions table automatically. A gap exists when one scenario's High multiplier is less than the adjacent scenario's Low multiplier while both scenarios have non-zero probability — for example, BTP High = 0.85 while AP Low = 0.90 (shown in Figure 4.2a). Cost multipliers in the gap range have zero probability under any scenario, producing a zero-density region in the combined distribution that can distort simulation results.

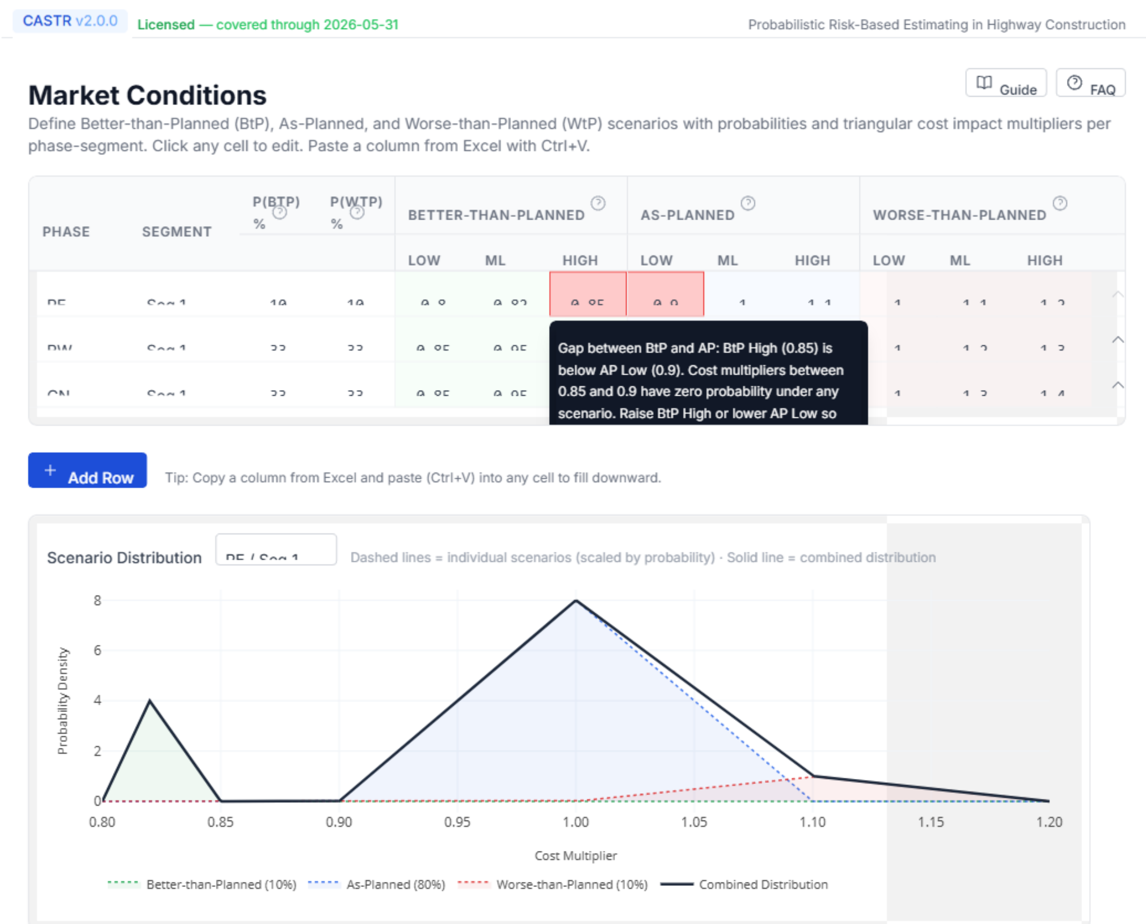

Figure 4.2a - Gap detection: BTP High (0.85) below AP Low (0.90). Both cells are flagged red; the combined distribution in the chart below shows a zero-density region between the BTP and AP triangles.

When a gap is detected, both cells that form the gap (BTP High and AP Low, or AP High and WTP Low) are given a red background. Hovering over a flagged cell displays a dark tooltip with the specific values involved and instructions on how to fix it (raise the lower field or lower the higher field so the scenario ranges overlap). The Scenario Distribution chart directly below the table shows the gap visually as a dip in the combined-distribution line.

Chapter 5: Risk Analysis

5.1 Risk Identification

New in v1.8: *This page now supports a Starter Library picker, a per-risk Mitigation Details panel, a Mitigation vs. Non-Mitigation filter, and shortened IN/MI/ME dependency codes. See Chapters 9 and 10 for details. New in v1.9: a SEL column with per-row checkboxes and a Delete Selected button for multi-row deletion (with an auto-checkpoint as the undo safety net); a help tooltip on the Details column header; a “(select Mitigate to access details)” hint on the Risk Response dropdown; the Description field in the Risk Details dialog renamed to Risk Statement with a new placeholder template; Residual P(RESID) is required when Risk Response is set to Mitigate (blank saves show a red banner); and save-time distribution validation (see Appendix B — Parameter Validation).*

The Risk Identification page is the primary interface for entering and managing individual risk events. Each risk represents a specific event or condition that could affect project cost, schedule, or both.

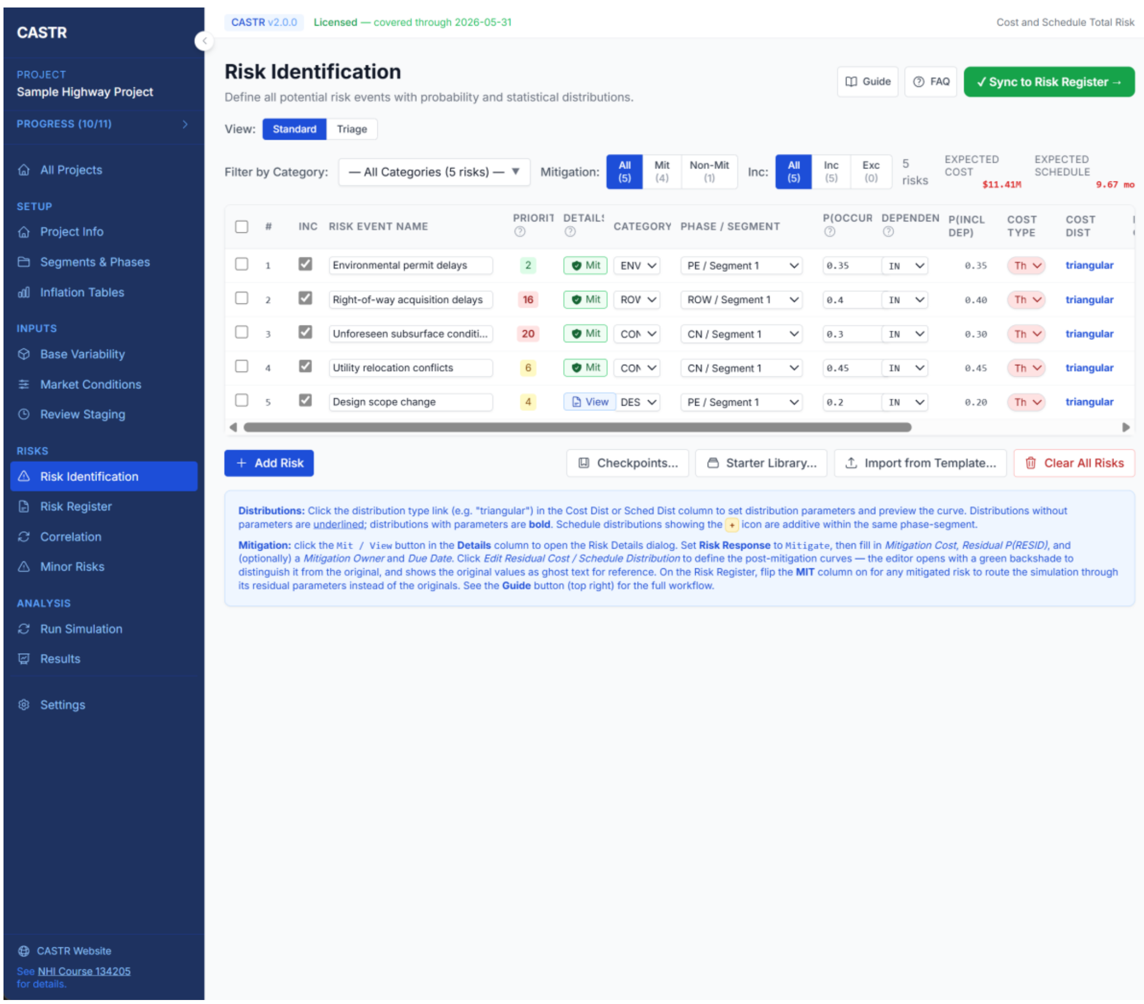

Figure 5.1 - Risk Identification page

The main risk entry table displays one row per risk event. Each risk has the following attributes:

- **Include:** A checkbox that controls whether the risk is included in the simulation. Unchecked risks are retained in the table but excluded from analysis.
- **Risk Event Name:** A descriptive name for the risk. Choose names that clearly communicate the nature of the risk event.
- **Category:** The risk category (e.g., ENV, CON, DES) as defined on the Segments and Phases page.
- **Phase/Segment:** The phase-segment combination that the risk applies to. A risk can only be assigned to one phase-segment.
- **Probability of Occurrence:** A value between 0 and 1 representing the likelihood that the risk event will occur. A value of 1.0 means the risk is certain to occur.
- **Dependency Type:** Defines how the risk relates to other risks. Independent risks are sampled individually. Mutually Inclusive risks always occur together when triggered. Mutually Exclusive risks cannot occur simultaneously.
- **Cost Distribution:** The probability distribution for cost impact. Click the distribution type link (e.g., "triangular") to open the Distribution Editor, which provides a visual preview of the PDF and CDF curves as you adjust parameters.
- **Schedule Distribution:** The probability distribution for schedule impact, configured the same way as the cost distribution.

Click the "View" button on any risk row to open the Details modal. This modal provides fields for documenting the risk description, planned response actions, risk owner, and current status. This information is carried through to the Risk Register for reporting purposes.

To import risks from an external source, click "Import from Template..." at the top of the page. This feature allows you to download a pre-configured Excel template, fill in risk data offline, and upload the completed template. Imported risks are validated and previewed before being added to the project.

Click "Sync to Risk Register" to push the current risk data to the Risk Register page, which provides additional reporting and documentation capabilities.

5.1.1 Standard View

The Standard view is the default layout for Risk Identification and shows the full set of risk attributes described above. Use this view for detailed risk entry, distribution editing, mitigation configuration, and any work that requires access to documentation fields.

Priority column (added in v1.12). The Standard view includes a Priority column immediately after Risk Event Name. It displays the Combined Priority Score computed in the Triage view (Probability score × max(Cost, Schedule)). Scores are color-coded: green (1–3), yellow (4–8), amber (9–15), red (16–25). Risks that have not yet been scored in Triage show a dash. The column is read-only here — edit scores on the Triage view.

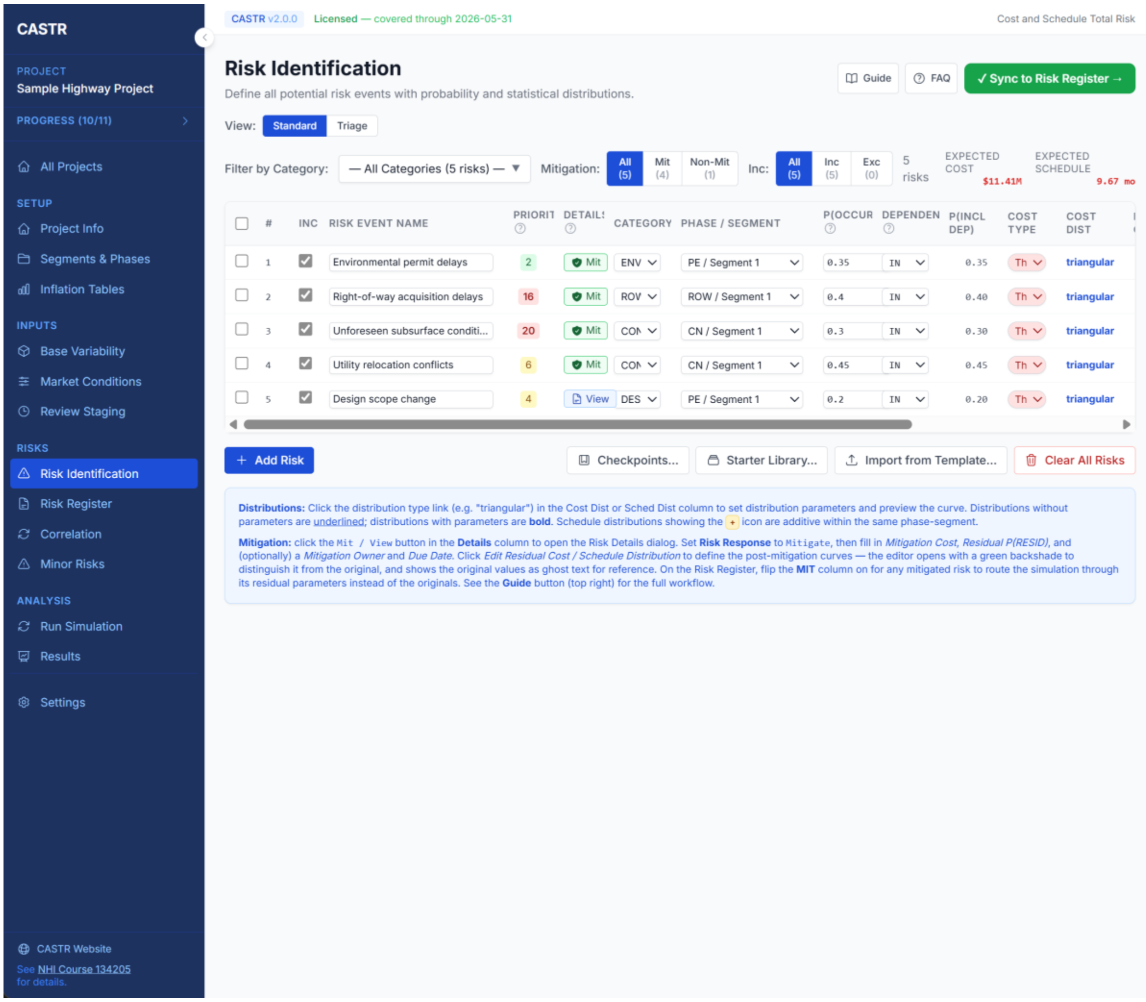

Figure 5.1a - Risk Identification (Standard view) with Priority column

5.1.2 Triage View (added in v1.9)

The Triage view is an alternative layout optimised for rapid, workshop-style scoring. Switch to it using the view toggle above the risk table. Each risk receives three 1–5 scores that summarise a qualitative assessment before committing to full probability distributions:

Prob (Probability) – the likelihood that the risk occurs. 1 = rare, 5 = near certain.

Cost – the magnitude of the cost impact if the risk occurs, as a relative severity against the phase-segment's Base Cost.

Sched – the magnitude of the schedule impact if the risk occurs.

Scores are entered using segmented 1–5 buttons (replacing the older dropdown). Click a number to set it; click the same number again to clear it. The buttons save immediately — no Enter key or blur event is required.

A Combined Priority score is computed automatically as Prob × max(Cost, Sched) and shown in the Priority column. The column header is not sorted by default; click it to toggle an ascending/descending sort. Multiple risks can be selected (checkbox in the # column) and deleted as a group.

Tip: Use the Quantify segmented filter (All / Quantify / Skip) above the table to divide risks into those that will graduate to full probability distributions and those that are better handled as Minor Risks. The Quantify flag is a simple toggle on each row.

Additional Triage filters (added in v1.10): two dropdowns beside the Quantify segmented control let you narrow the table further. Category filters to risks tagged with a specific category (or "Uncategorized"); Phase / Segment filters to risks assigned to a specific phase-segment (or "Project Wide"). When any Triage filter is active a "Clear filters" link appears beside the controls - one click resets Quantify, Category, and Phase/Segment at once. The Triage table also gained a sortable "Risk Event Name" column header - click it to sort A->Z or Z->A, composable with the existing Priority / Category / Phase-Segment sort options.

5.1.3 Heat Maps (added in v1.9)

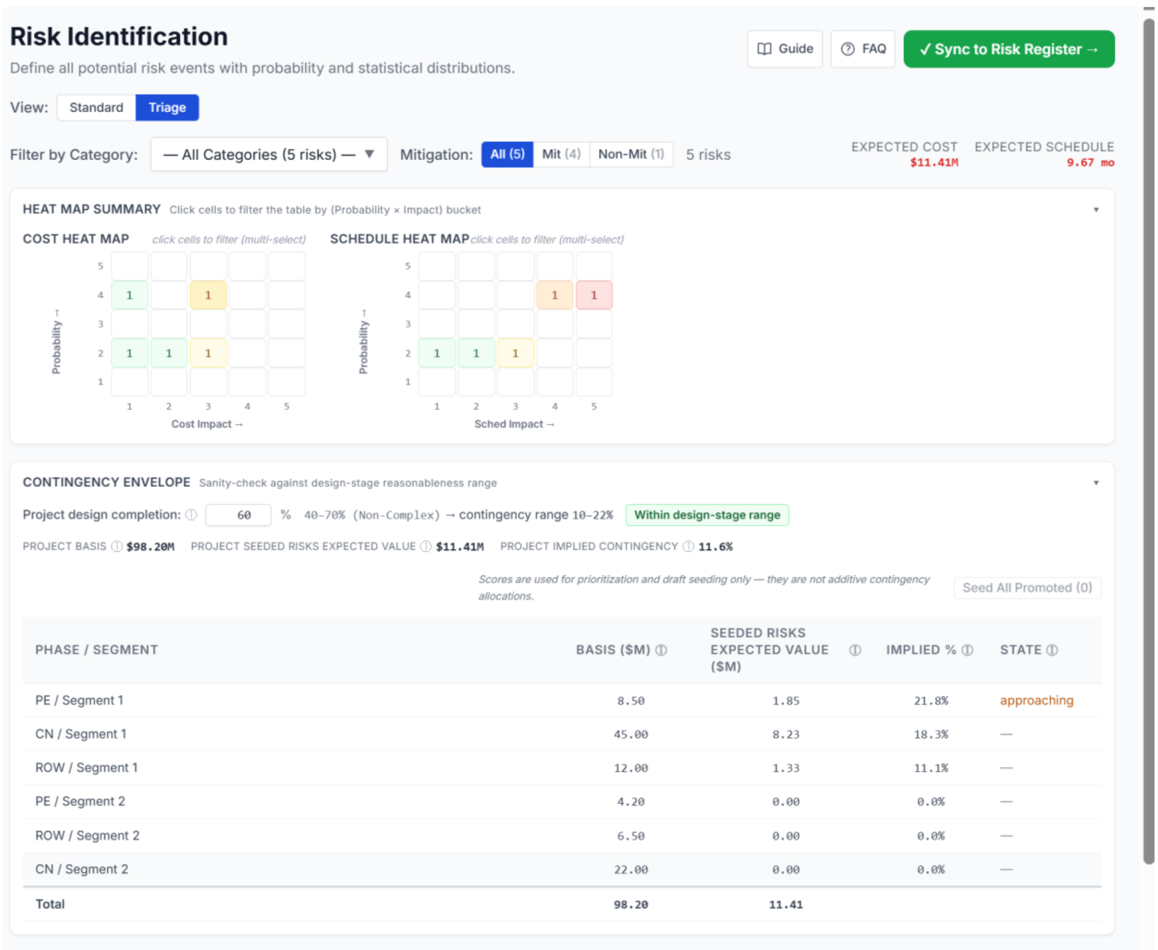

Figure 5.1b - Cost and Schedule Heat Maps (Triage View)

Click any cell to filter the Triage table to only risks falling in that cell. Click the cell a second time to deselect it. Multiple cells can be selected at once, in either map or across both: the Triage table shows the union of risks matching any selected cell.

Tip: to isolate a group of related cells (e.g. every Cost cell in the top-right quadrant), simply click each one in turn. Use the Clear heat-map filter link that appears below the maps to remove all selections at once.

Warning: heat-map filters compose with the category, mitigation-status, and Quantify filters. If the table is unexpectedly empty, check that another filter isn't restricting the set — the quickest way to reset is to click the filter chips above the table.

5.1.4 Seeding Distributions from Triage Scores (added in v1.9)

The Seed button on each Triage row generates starter Triangular cost and schedule distributions from the risk's 1–5 scores. This turns a qualitative assessment into concrete low/mode/high parameters in one click, so risks can graduate from triage to quantification without typing numeric estimates by hand. A Seed-All button is also available for bulk seeding every Quantify-flagged risk that currently has no distribution.

How the values are produced:

Cost: the risk's phase-segment Base Cost (as defined in the Glossary — the deterministic current-year cost with the Traditional Contingency line already removed; this value is shown as the "Base Cost" total on the Base Variability page) is multiplied by the low/mode/high fractions from the Cost scoring table for that cost_score. The result is the raw starter triangle.

Schedule: the low/mode/high weeks from the Schedule scoring table for that schedule_score are converted to months (÷ 4.3482 weeks/month). Schedule impacts are absolute durations, so no basis is required.

Calibration. CASTR scales the raw values by a phase-segment size factor so that risks on unusually small or large phase-segments do not produce unrealistic distributions. The factor equals clamp(this_phase_segment_basis ÷ project_mean_phase_segment_basis, 0.75, 1.25). A risk on an average-sized phase-segment is unchanged (factor = 1.00). A risk on a very small phase-segment is floored at 0.75× so its band is not driven unrealistically low; a risk on a very large phase-segment is capped at 1.25× so it is not inflated above the intended sensitivity range. Note: when the calibration factor reaches the 0.75 or 1.25 boundary, the persisted distribution is materially scaled relative to your scoring-table value. Confirm this in the Notes section of the Seed dialog before clicking OK.

Confirmation dialog: when a risk already has a cost or schedule distribution set, clicking Seed opens a confirmation dialog showing the current distribution (with its type — Triangular, BetaPERT, Uniform, Normal, or Constant — and parameters), the proposed calibrated starter values that would be written, and a Notes section that explains any calibration factor applied and any clamp that fired. Only the proposed (calibrated) values are written if you click OK. Cancel leaves the existing distribution intact.

Tip: seeded distributions are intended as a starting point. Open the Distribution Editor from the risk's row (or switch back to Standard view) to refine the low/mode/high values or change the distribution type. Seeding does not lock the risk — subsequent edits are fully permitted.

Warning: Seed uses the Cost and Schedule scoring tables currently in effect. If you have edited these tables on the Settings → Scoring Tables page, the starter values will reflect your edits, not the built-in defaults. Already-seeded distributions are NOT updated retroactively when the tables change — only subsequent Seed operations use the new values.

Probability seeding from Prob score (added in v2.2). When a risk's P(occur) cell is empty, Seed Distributions also writes a P(occur) value derived from the Triage Probability score (1–5) via the Probability Bands table on Settings → Scoring Tables. The factory defaults are 1=0.10, 2=0.30, 3=0.50, 4=0.70, 5=0.90 — i.e., a Prob score of 3 seeds P(occur)=0.50 unless the agency has customized the bands. User-typed P(occur) values are never overwritten by Seed; the seeding only fills blank cells.

Yellow shield badge — P(occur) cells populated by Seed (rather than typed by a user) display a small yellow shield badge as a “seeded — please confirm” nudge to the workshop team. The badge clears the moment the cell is edited. Two helpers make the seeded value predictable in a workshop setting: a “?” help button next to the Triage Prob column header opens a panel showing the live 1–5 → P(occur) mapping, and the 1–5 score chips themselves show the mapped P(occur) on hover so a facilitator can preview what each click will produce.

5.1.5 Triage Contingency Envelope Panel (added in v1.9; SME columns moved out in v2.5)

A collapsible Contingency Envelope panel sits above the Triage table showing the Project Implied Contingency against an industry-standard envelope for the project's design completion stage. The envelope is drawn from the Design Completion bands — by default, a 1–10 % design (Concept) project has an envelope of 30–75 % total contingency, narrowing to 7–20 % at 70–100 % design (Final). The project's Design Completion percentage is set on the Project Information page. The Triage card has six sortable per-phase-segment columns (Phase / Segment, Basis, Seeded EV, Seeded Implied %, Risks (Active / Reviewed / Inactive), and State); the SME-Reviewed columns moved to the Standard view envelope card (§5.1.6) and the Risk Register envelope card (§5.2) in v2.5. A project-level SME-Reviewed Implied chip is retained in the Triage card's header so the workshop can see overall ratification progress without leaving Triage.

Three envelope states are possible:

Within envelope – the implied contingency fits inside the band. A green check is displayed.

Approaching envelope – the implied contingency is above 85 % of the band's high end. An amber chip appears as a soft warning that additional seeding may push the project over the envelope.

Exceeded envelope – the implied contingency is above 105 % of the band's high end. A red chip appears, and the bulk Seed-All action requires an additional confirmation before it will proceed.

Tip: the envelope is a sanity check, not a hard rule. An exceeded envelope can indicate that (a) your Cost scoring table is tuned too hot for this project type, (b) several risks are incorrectly flagged for Quantify when they belong on the Minor Risks list, or (c) the project genuinely carries unusually high risk for its design stage and the envelope itself is inappropriate. Addressing (a) or the envelope itself is a Settings → Scoring Tables edit; addressing (b) is a per-risk triage decision.

Warning: the envelope panel is a cost-only check. Editing the Schedule scoring table does not affect the envelope; editing the Cost scoring table OR the Design Completion bands does.

State column tooltip (added in v1.10): the State column in the Contingency Envelope panel now includes a help icon explaining the possible values. OK means the Implied Contingency falls within the design-stage reasonableness range. LOW flags a phase-segment whose implied contingency is below the expected range (potentially under-seeded risks). HIGH flags one above the range (potentially over-seeded, or a signal of over-exposure). Em dash ("--") appears when no Design Completion % has been entered. Use the State column to steer the workshop toward the phase-segments that most need attention rather than reviewing every risk equally.

5.1.6 Standard View Contingency Envelope (SME-Reviewed; added in v2.5)

The Standard view shows its own Contingency Envelope card directly above the risk table. Where the Triage card (§5.1.5) is the seed-stage view (every priced risk counts toward the totals), the Standard card is the post-workshop, defensible view — it counts only risks the SMEs have explicitly ratified. A risk qualifies as SME-Reviewed when both its cost-distribution unconfirmed flag AND its probability unconfirmed flag have been cleared (open the Distribution Editor and click Save; press Enter on P(occur)) — see §12.1.1.

The card is collapsible. Project-level totals shown in the header: Project Basis, Project SME-Reviewed Risks Expected Value, Project SME-Reviewed Implied %, and a Within / Approaching / Exceeded chip vs. the design-stage range. The per-phase-segment table has six columns: Phase / Segment, Basis ($M), SME-Reviewed Risks Expected Value ($M), SME-Reviewed Implied %, Risks (Active / Reviewed / Inactive), and State. Cells render an em dash (—) where nothing has been reviewed yet for that phase-segment. The State header has a help icon explaining the possible values (⚠ exceeded, approaching, —).

The card is read-only here. Edits to the project's Design Completion % are made on Risk Identification → Triage to keep a single input surface for that project field. To track workshop progress, compare the Standard card's SME-Reviewed Implied % against the Triage card's Seeded Implied % — when the two converge, the contingency is fully ratified.

5.2 Risk Register

New in v1.8: The Risk Register now includes an MIT column and related mitigation columns that appear automatically when at least one risk on the project has Risk Response = Mitigate. Sync now preserves 1:1 numbering between Risk Identification and the Register. See Chapter 10. New in v1.9: the Description column is renamed to Risk Statement; the page header is reorganized (Guide on the title line at far right, Re-sync / Print / CSV / Excel on the subtitle line at far right, Checkpoints button moved below the table); a workshop-visible row hover highlight (strong brand-blue fill with a 4px left accent bar) makes the active row easy to track during SME workshops; and the MIT toggle now has a dedicated explanation in the Guide covering activation, eligibility, workflow, and re-sync protection.

The Risk Register provides a comprehensive, report-oriented view of all risks that have been synced from the Risk Identification page. It is designed for documentation, review meetings, and formal risk reporting.

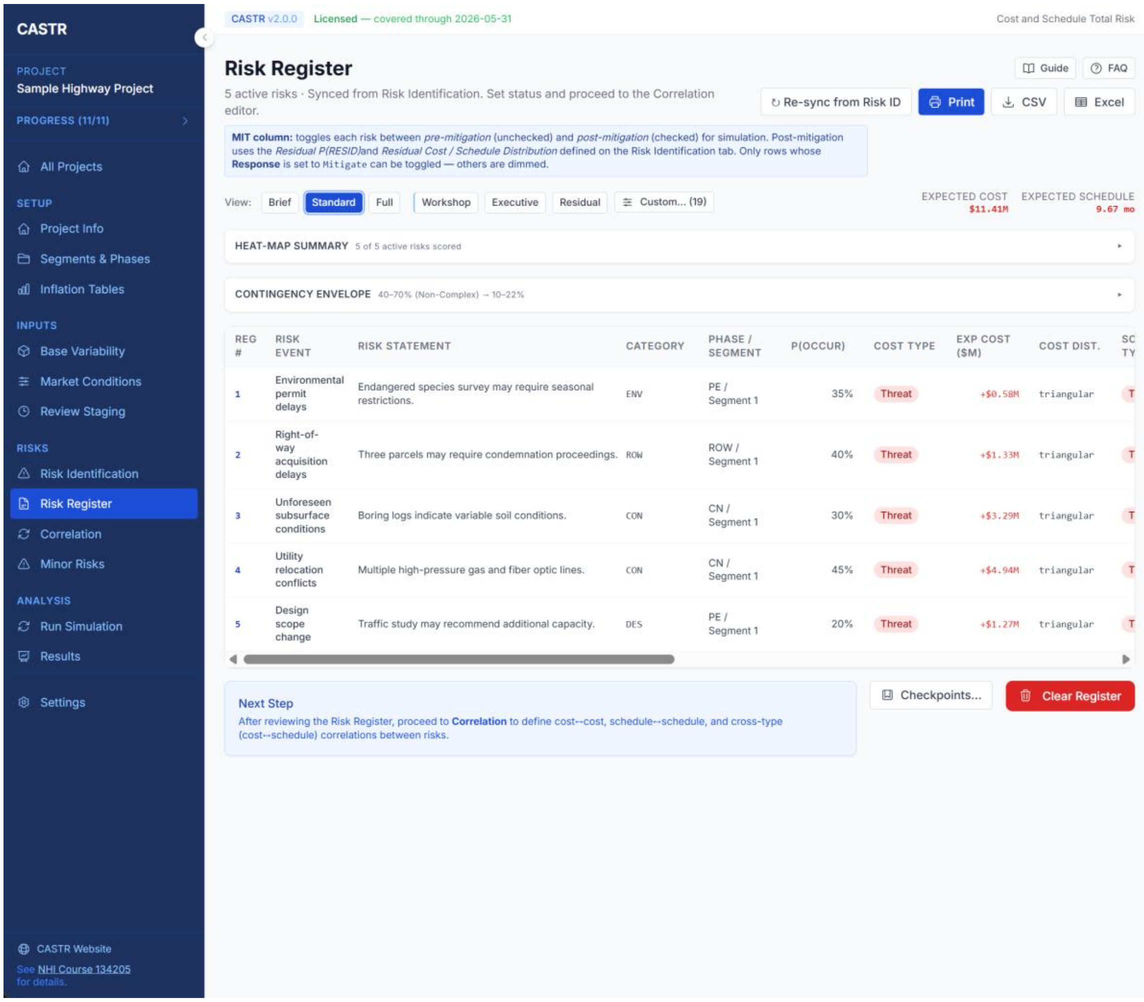

Figure 5.2 - Risk Register page

The Risk Register offers multiple view presets to control which columns are displayed:

- **Brief:** Shows only the essential columns: risk name, category, probability, and cost/schedule impact.
- **Standard:** Adds status, owner, and response action columns to the Brief view.
- **Full:** Displays all available columns including detailed documentation fields.
- **Custom:** Allows you to select exactly which columns to display.

Column groups include Identification (name, category, phase-segment), Probability and Impact (occurrence probability, cost and schedule distributions), Status and Ownership (risk status, owner, response actions), and Documentation (description, notes, last updated).

You can export the Risk Register to CSV or Excel format using the buttons at the top of the page. A print-optimized layout is also available for generating hard-copy reports.

Individual rows can be locked to prevent them from being overwritten when you re-sync from Risk Identification. This is useful when you have added documentation or status updates directly in the Risk Register that you do not want to lose. You can also set the risk Status to Active, Inactive, or Closed to track the lifecycle of each risk.

Contingency Envelope Card (revised in v2.5 — SME-Reviewed): The Risk Register includes a collapsible Contingency Envelope card (above the risk table) that shows the defensible / sign-off-ready picture of the project. Only risks the workshop has explicitly ratified contribute to its totals — see §12.1.1 for what counts as "SME-Reviewed". Expanding the card reveals the project-level summary metrics — Project Basis, Project SME-Reviewed Risks Expected Value, and Project SME-Reviewed Implied % — followed by a per-phase-segment breakdown with six columns: Phase / Segment, Basis ($M), SME-Reviewed Risks Expected Value ($M), SME-Reviewed Implied %, Risks (Active / Reviewed / Inactive), and State. Cells render an em dash (—) where nothing has been reviewed yet. Hover (or click) the small question-mark icons to see inline definitions for each metric. The card is informational only and read-only: state is driven by values entered on other pages (Base Cost, risk probabilities / distributions, and the Design Completion % on Risk Identification → Triage). Compare this card's SME-Reviewed Implied % against the Seeded Implied % on Risk Identification → Triage — when the two converge, the contingency is fully ratified.

5.3 Correlation Editor

The Correlation Editor allows you to define statistical correlations between risk variables. Correlations describe how the sampled values of two risks move together during simulation. Positive correlation means the risks tend to be high together and low together; negative correlation means one tends to be high when the other is low. For example, if material costs increase, labor costs may also increase due to broader inflationary pressures.

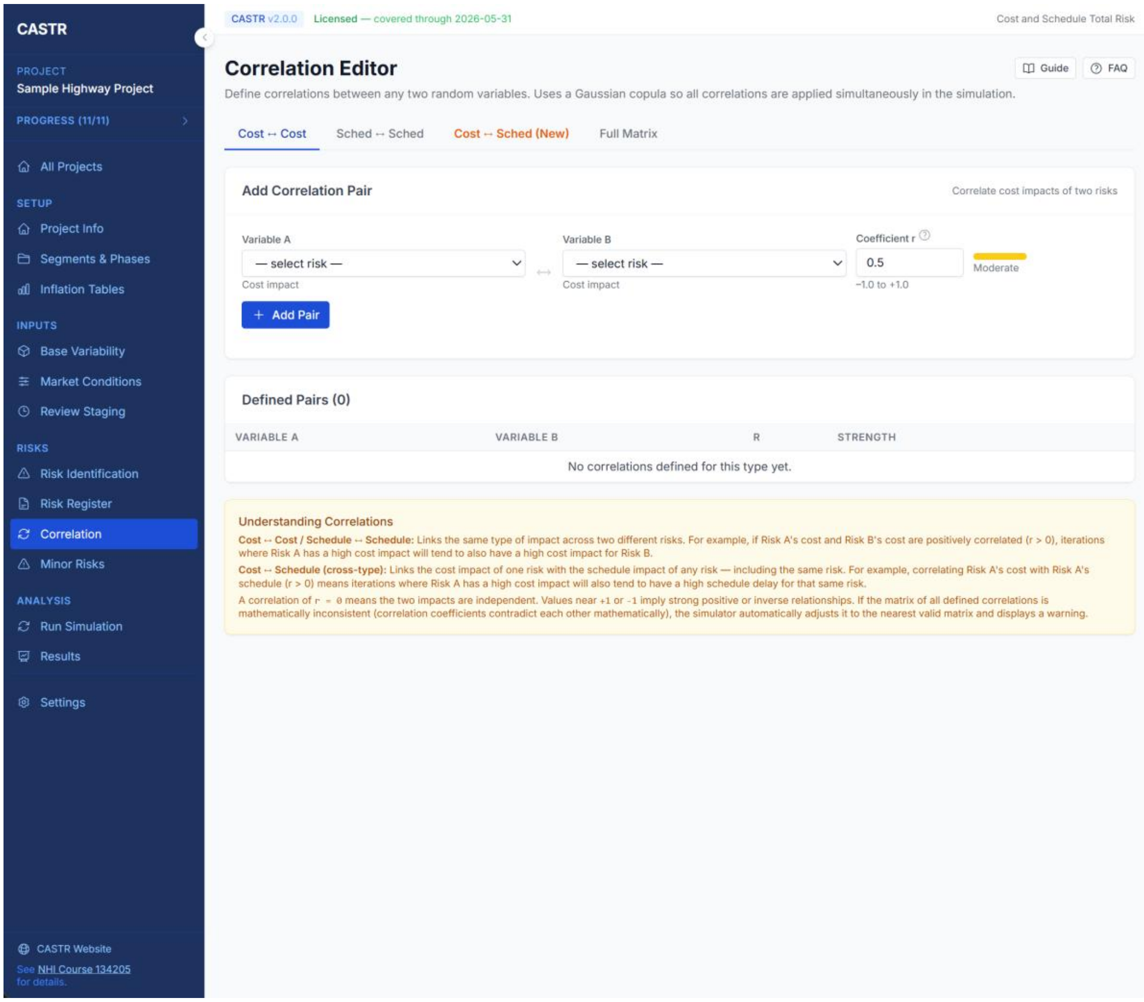

Figure 5.3 - Correlation Editor page

The editor is organized into four tabs for different correlation types:

- **Cost-Cost:** Correlations between the cost impacts of different risk variables.
- **Schedule-Schedule:** Correlations between the schedule impacts of different risk variables.
- **Cost-Schedule:** Cross-type correlations between the cost impact of one variable and the schedule impact of another.
- **Full Matrix:** A comprehensive view of all correlations in matrix format.

To define a correlation, select an anchor risk in Variable A and one or more target risks in Variable B (a multi-select picker, added in v2.2.1), then set the correlation coefficient r, which ranges from -1 to +1. A value of +1 indicates perfect positive correlation (both variables move in the same direction), -1 indicates perfect negative correlation, and 0 indicates no correlation.

CASTR uses a Gaussian copula to apply all correlations simultaneously during simulation. This mathematical method ensures that the correlations between all pairs of variables are applied consistently. If the resulting correlation matrix is not positive-semi-definite (PSD) — which can happen with certain combinations of correlation values — the simulator automatically adjusts the matrix to the nearest valid PSD matrix using the Higham (2002) algorithm and proceeds with the simulation. See the PSD Validity note below.

One-to-Many Entry (added in v2.2.1). On each of the three editable tabs (Cost ↔ Cost, Sched ↔ Sched, and the Cost ↔ Sched cross tab), Variable B is a multi-select picker rather than a single dropdown. The user picks one anchor risk on the left, ticks any number of target risks on the right, and presses Add — CASTR writes one pairwise CorrelationEntry per (anchor, target) pair, all with the same coefficient. The button label changes to reflect the count (“Add Pair” for one target, “Add 5 Pairs” for five). Selected targets appear as removable chips below the form so the user can scan exactly what will be saved before pressing Add.

The picker has a search box (filters by risk number or name), a “Select all visible” shortcut that respects the current filter, and a Clear button. After save, the anchor and coefficient are retained so the user can iterate quickly; only the target list resets. This is the natural way to express workshop statements like “Schedule delays on Risk 1 correlate with the cost of Risks 1, 2, 4, 7, and 8 at r = 0.5” — one anchor, five targets, one coefficient, one click.

Conflict handling. CASTR detects when a target pair already exists in either direction (an existing A↔B entry collides with a draft B↔A pair too). Pairs that already match at the same r value are silently skipped. Pairs whose existing r differs from the draft value prompt a confirmation dialog: OK overwrites the existing entries with the new value; Cancel skips just the colliding ones (new pairs are still added). A small amber banner inside the form tracks the number of pairs that will be overwritten so the user sees the impact before pressing Add.

What it doesn't do. The fan-out shortcut produces one anchor → many targets — it does NOT produce a clique where every selected target also correlates with every other selected target. To build a true clique at one r, iterate the fan once per anchor (anchor R1 → targets R2-R5; anchor R2 → targets R3-R5; and so on). High-density cliques tend to violate positive semi-definiteness, so the engine's automatic Higham correction may flag the resulting matrix — see the PSD Validity note below for guidance on softening strong correlations.

PSD Validity. PSD = Positive-Semi-Definite. When you enter cost-cost, schedule-schedule, and cost-schedule cross-pair correlations, the resulting correlation matrix must be positive-semi-definite for the Gaussian copula simulation to work mathematically. A matrix fails the PSD test when entered correlations are internally inconsistent — for example, A is highly correlated with both B and C while B and C are uncorrelated with each other. CASTR detects this automatically and repairs the matrix to the nearest valid PSD matrix using the Higham algorithm, then reports an achieved-vs-requested correlation comparison so you can see what was adjusted. To soften strong correlations in a dense clique, reduce r values on pairs until the comparison report shows negligible adjustments — values in the 0.4–0.6 range rarely trigger corrections, while dense cliques with r ≥ 0.8 commonly do.

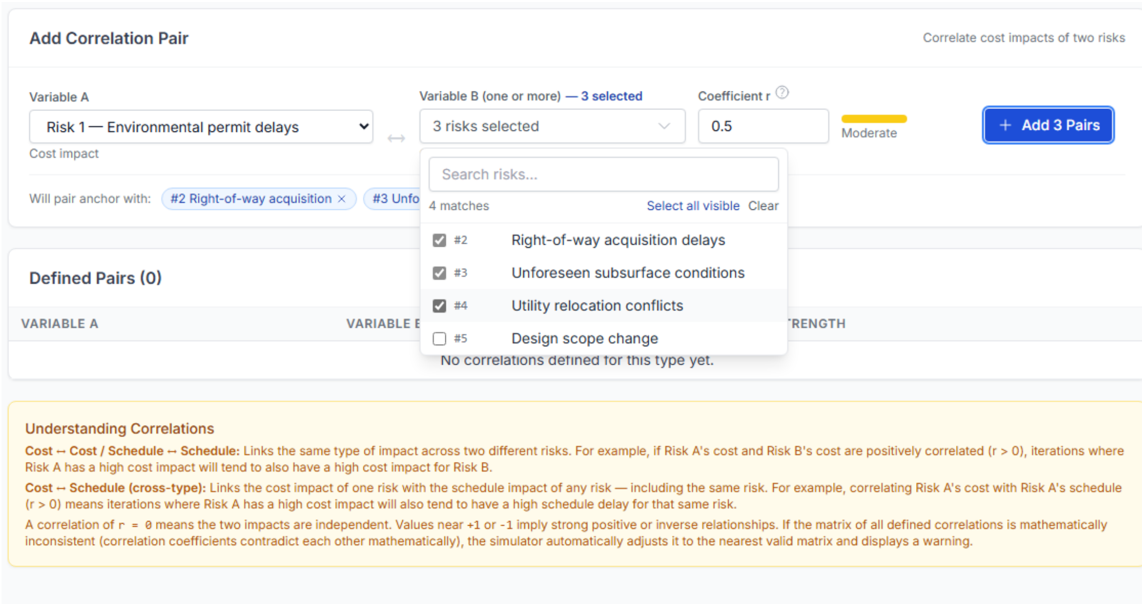

Figure 5.3a — One-to-Many correlation entry. Variable A is a single anchor; Variable B is a multi-select picker. The chip row below the form previews exactly which (anchor, target) pairs will be written on Save. The Add button label reflects the count (here: "Add 3 Pairs").

5.4 Minor / Aggregate Risks

The Minor Risks page provides a simplified entry method for low-impact risks that do not warrant full probabilistic modeling. These are risks that individually have small impacts but collectively can contribute meaningfully to the total project risk profile.

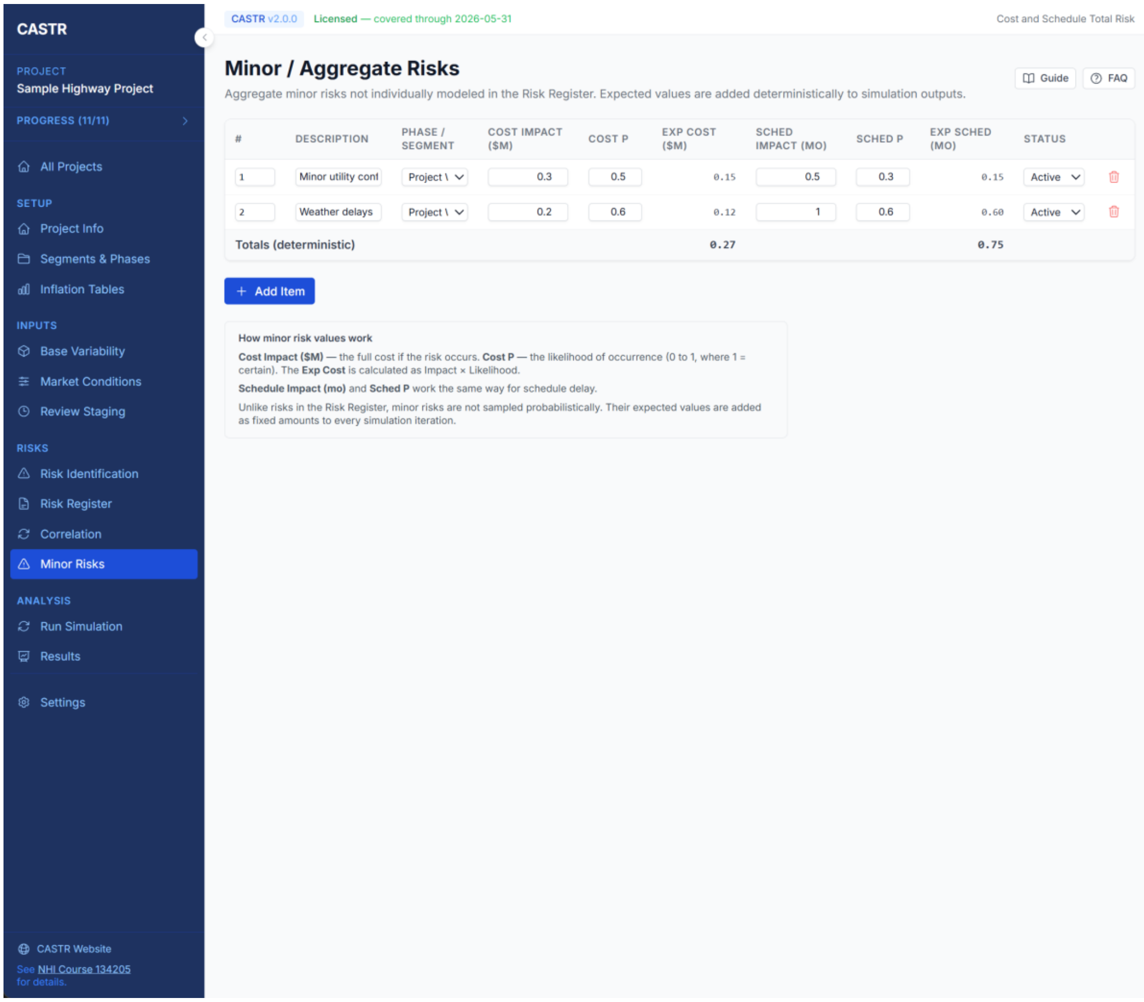

Figure 5.4 - Minor / Aggregate Risks page

For each minor risk, you enter the following fields:

- **Description:** A brief description of the risk event.
- **Phase/Segment:** The phase-segment combination affected by this risk.
- **Cost Impact ($M):** The estimated cost impact in millions of dollars if the risk occurs.
- **Cost Probability:** The probability (0 to 1) that the cost impact will occur.
- **Schedule Impact (months):** The estimated schedule delay in months if the risk occurs.
- **Schedule Probability:** The probability (0 to 1) that the schedule impact will occur.

Expected values are automatically calculated as the product of impact and probability for both cost and schedule. Unlike risks entered on the Risk Identification page, minor risks are not sampled probabilistically during simulation. Instead, their expected values are added deterministically to every simulation iteration. This simplifies the modeling of numerous small risks while still capturing their aggregate contribution to the project cost and schedule.

5.5 Review Staging

The Review Staging page allows you to review the phase-segment staging and assign predecessor and successor linkages to reflect dependencies so the schedule reflects the real order of project delivery.

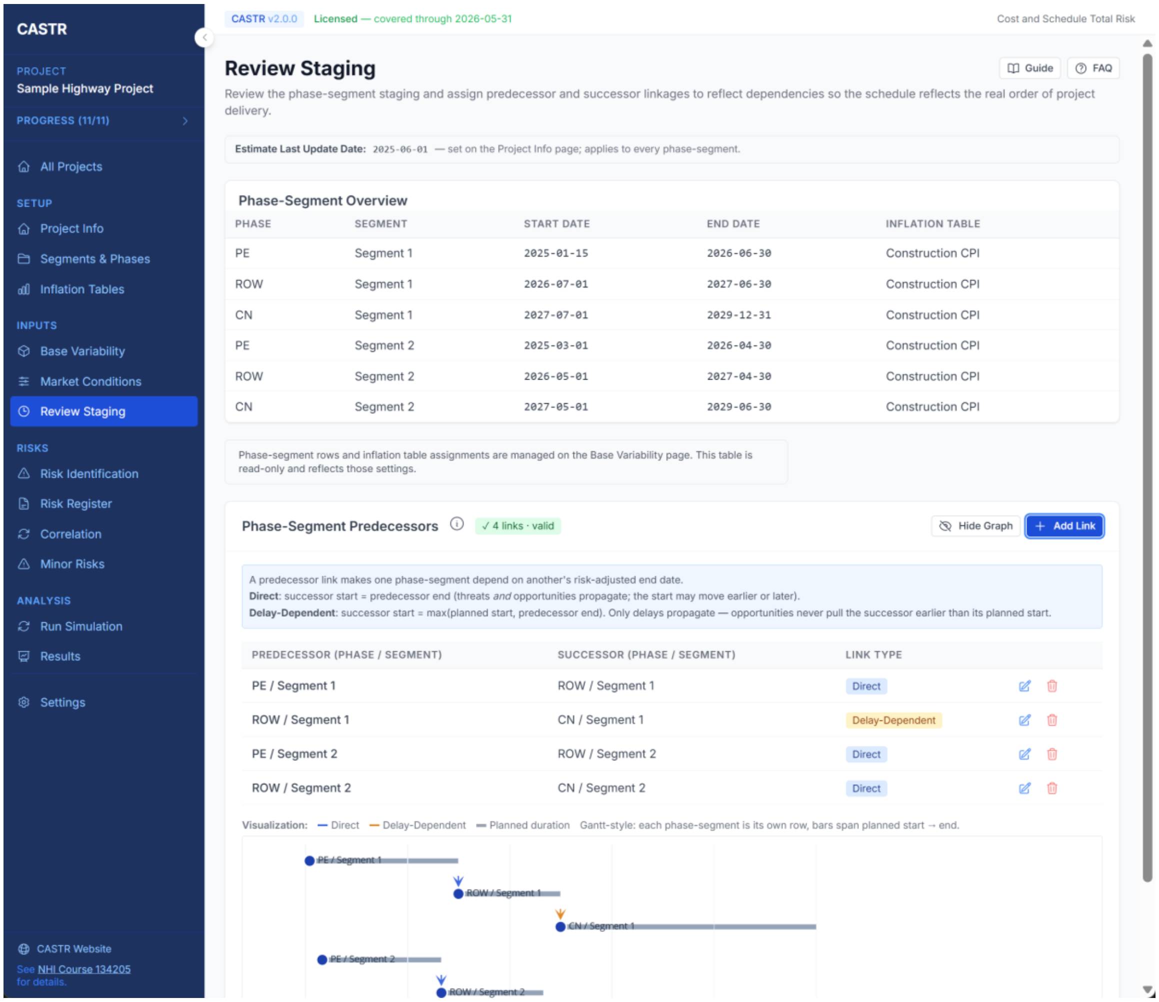

Figure 5.5 - Review Staging page

Phase-Segment Overview: The top section displays a read-only summary table showing the start and end dates, and the assigned inflation table for each phase-segment. These values are managed on the Base Variability and Inflation Tables pages respectively. This table serves as a convenient reference when setting up predecessor relationships.

Phase-Segment Predecessors: The lower section allows you to define execution dependencies between phase-segments. Predecessors determine the order in which phase-segments are executed and how schedule changes in one phase-segment propagate to others.

Two link types are available for predecessor relationships:

- **Direct (D):** The successor phase-segment starts immediately when the predecessor ends. Both delays and schedule improvements (opportunities) in the predecessor propagate to the successor.
- **Delay-Dependent (DD):** Only delays in the predecessor propagate to the successor. If the predecessor finishes early (an opportunity), the successor does not start earlier than

originally planned. This is useful when a successor cannot practically start earlier even if the predecessor completes ahead of schedule.

A Gantt-style visualization is displayed to help you understand the dependency structure and verify that the relationships are configured correctly.

Chapter 6: Running the Simulation

6.1 Simulation Settings

***New in v1.8:** A new Run Mode setting lets you run the simulation in Pre-Mitigation, Post-Mitigation, or Compare-Both mode. See Chapter 10. New in v1.9: market-condition and risk-event rows that are still being filled in (blank phase/segment, blank multipliers, or incomplete triangular triples) are safely skipped by the engine rather than aborting the run.*

The Simulation page allows you to configure the Monte Carlo simulation parameters before running the analysis. Proper configuration of these settings ensures that the simulation produces reliable and stable results.

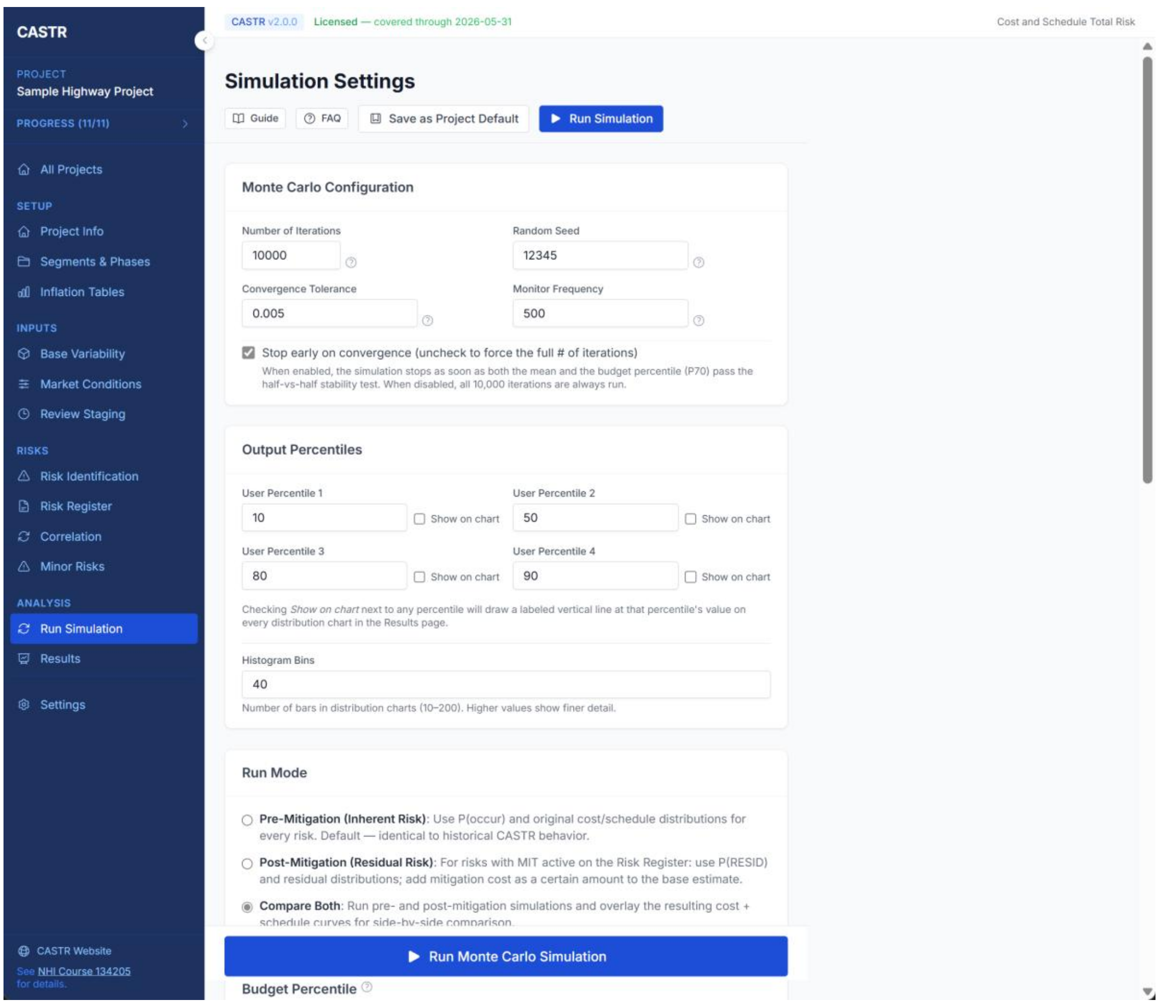

Figure 6.1 - Simulation Settings page

The following Monte Carlo configuration parameters are available:

- **Number of Iterations:** The total number of complete runs of the model using one random draw for each uncertain input per run. The minimum is 100 and the maximum is 100,000 iterations. The default is 10,000, which is sufficient for most projects. Increase this value for projects with many risk variables or when higher statistical precision is required.

- **Random Seed:** An optional seed value for the random number generator. Setting a seed ensures that the simulation produces identical results each time it is run with the same inputs, which is useful for reproducibility and auditing.
- **Convergence Tolerance:** The tolerance for convergence detection, as a fraction (e.g., 0.005 = 0.5%). The simulation checks whether P10 and P90 have stabilized. If the change between the first and second half of iterations is below this tolerance, the simulation stops early. The default of 0.005 (0.5%) is a balance between speed and stability for typical highway projects. It bounds the relative change in P10 and P90 between halves of the run; it does not by itself bound the absolute estimation error of any single percentile.
- **Monitor Frequency:** Monitor Frequency sets how many iterations run before the simulation pauses to update progress, check for cancellation, and, if enabled, test for convergence. Lower values increase responsiveness but add slight overhead. Higher values reduce overhead but make updates, cancellation, and convergence checks less frequent. The default value, 500, is a good balance.

Output Percentiles: Configure which percentile values are tracked and displayed in the results. Common percentiles include P10, P50, P80, and P90. The P50 (median) represents the value where half of all simulation outcomes fall below, while P80 means 80% of outcomes are at or below this value.

Budget Percentile: Choose how the headline budget number is computed from the per-iteration outputs. Three methods are available: Percentile of Sum (PoS) — the chosen percentile of the total project cost distribution; Sum of Segment Percentiles (SoSP) — the chosen percentile of each segment's distribution, summed; and Sum of Phase-Segment Percentiles (SoPSP) — the same idea at the finer phase-segment grain. PoS is generally recommended because it accounts for diversification across segments; SoSP and SoPSP are more conservative. The Reconstructed Distribution view on the Distribution Charts tab (see §7.2) reconciles the chart with whichever sum-of-percentiles method you chose.

Histogram Bins: Controls the number of bins in the distribution charts. Values range from 10 to 200. More bins provide finer granularity but may appear noisy with fewer iterations.

Click "Save as Project Default" to save the current settings so they are automatically loaded the next time you run a simulation for this project.

Inflation Sensitivity Scenarios (added in v2.2). Below the main settings is a panel where you can configure up to two ADDITIONAL inflation scenarios alongside the always-on baseline (delta = 0). Each enabled scenario takes a delta in PERCENTAGE POINTS that is added to every annual rate in every inflation table on the project, plus an optional human-readable label (e.g. “Recession”, “Overheated”). Deltas are capped at ±10 pp.

When at least one extra scenario is enabled, CASTR runs the simulation N times (1 baseline + up to 2 deltas), all sharing the same random seed so the only difference between runs is the inflation assumption. The Results page then shows a scenario chip selector at the top, a side-by-side comparison table on the CASTR Contingency card, and a small Sensitivity Curve chart (see §7.5).

Cost: thanks to internal random-draw caching, the typical 3-scenario run completes in roughly 2.5× the time of a single-scenario run (rather than 3×). The cache lets every scenario after the first skip correlated-uniform sampling, ICDF transforms, Bernoulli risk-occurrence draws, and market-multiplier sampling. Schedule, completion-date, and Current-Year cost outputs are scenario-invariant by construction (only the CY → YOE inflation factor differs).

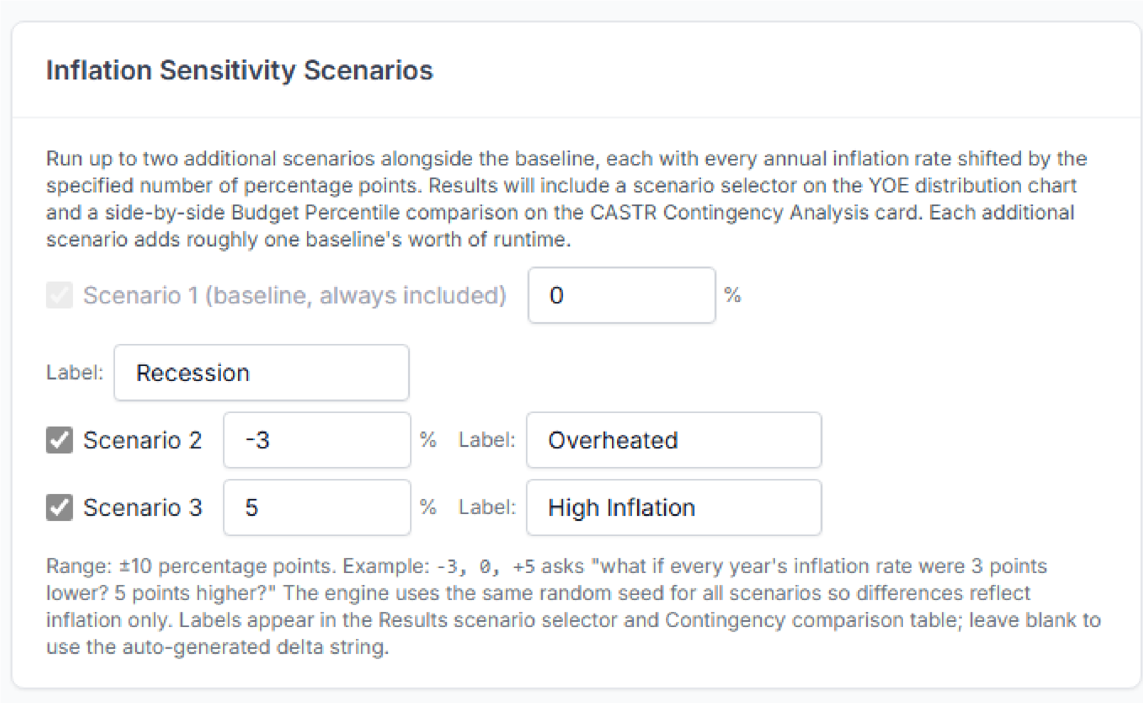

Figure 6.2 — Inflation Sensitivity Scenarios panel on the Simulation page. Scenario 1 is the always-on baseline (delta = 0). Up to two additional scenarios can be enabled with their own deltas (in percentage points) and optional labels. Empty labels fall back to the auto-generated delta string (e.g. "−3 pp").

6.2 Running and Monitoring

To start the simulation, click the "Run Monte Carlo Simulation" button. A progress bar will appear showing the current iteration count, total iterations, and elapsed time. A Cancel button is available if you need to stop the simulation before it completes.

The simulation runs in a background thread, so the user interface remains responsive during execution. You can navigate to other pages while the simulation is running, although it is recommended to wait for completion before modifying project data.

When the simulation completes, results are automatically displayed on the Results page. If convergence was enabled and the simulation stopped early, a message indicates the iteration at which convergence was achieved.

6.3 Pre-Run Validation

Before starting the simulation, CASTR performs a validation check on your project data to ensure all required information is present and consistent. If any errors are found, the simulation will not start and error messages will be displayed describing the issues that need to be resolved.

Common validation errors include:

- No active phase-segment combinations defined on the Segments and Phases page.
- Missing Estimate Last Update Date on the Project Information page.
- No components enabled for simulation (all risks unchecked, no base variability data).
- Incomplete required data fields in the Base Variability table (missing dates or cost estimates).

- Invalid distribution parameters for risk entries (e.g., minimum value greater than maximum).

Resolve all reported errors and click Run again to start the simulation.

Chapter 7: Understanding Results

7.1 Summary Statistics

The Results Summary page provides a high-level overview of the simulation outcomes, including key statistical measures and the recommended risk-based contingency.

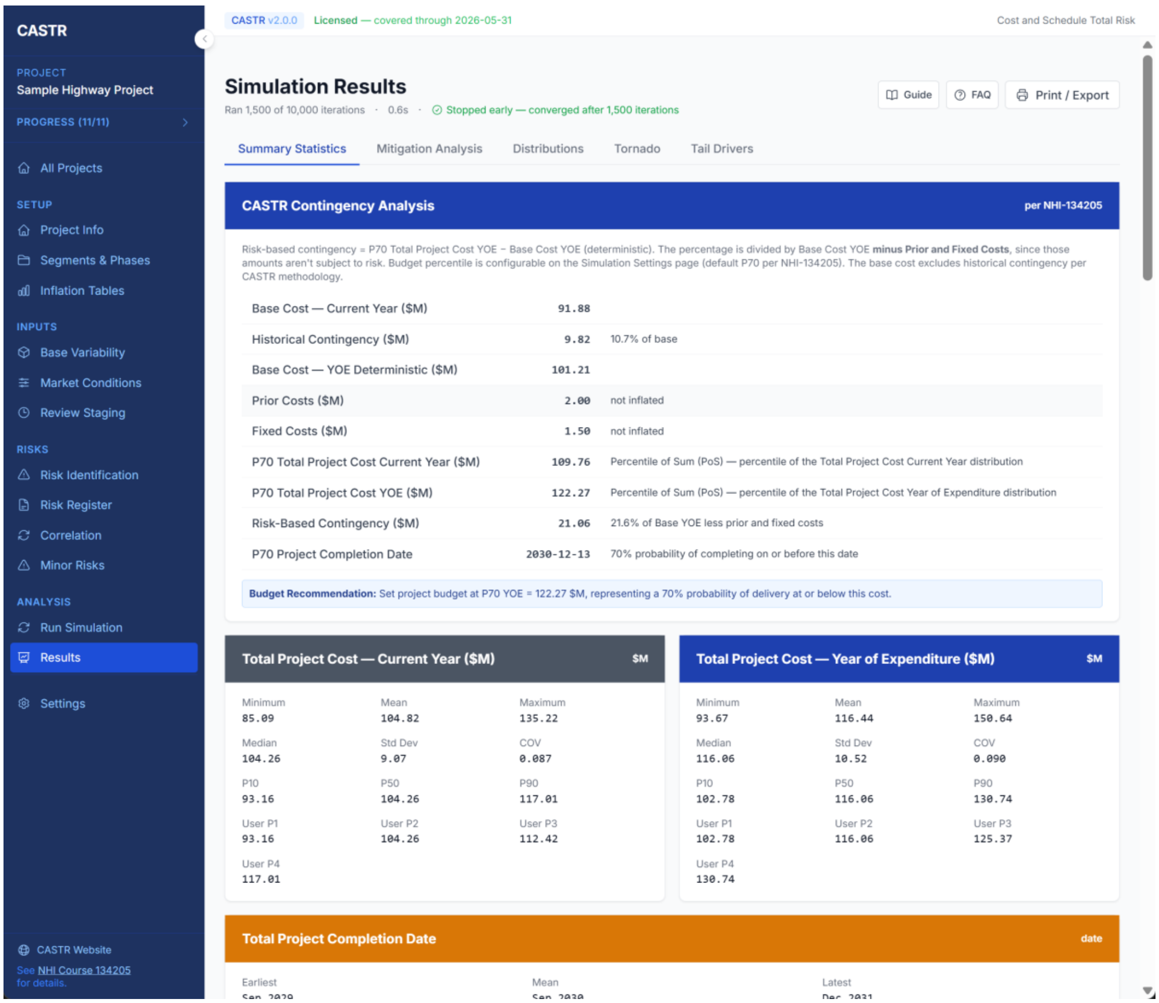

Figure 7.1 - Results Summary Statistics

CASTR Contingency Analysis Card: This card displays the base cost (the deterministic estimate without contingency), the historical contingency amount, prior and fixed costs, the budget recommendation at the percentile selected on the Simulation page (P70 by default), and the calculated risk-based contingency. The risk-based contingency is the difference between the budget percentile cost and the deterministic base cost, representing the additional funds needed to achieve the specified confidence level.

Two statistics cards are displayed below the contingency analysis:

- **Total Project Cost (Current Year):** Shows the distribution statistics for total cost in today's dollars, including Minimum, Mean, Maximum, Median, Standard Deviation, Coefficient of Variation (COV), and the user-selected percentile values.

- **Total Project Cost (Year of Expenditure):** Shows the same statistics but with costs adjusted for inflation to the years in which expenditures will actually occur.

7.2 Distribution Charts

The Distribution Charts page displays interactive Plotly charts that visualize the full probability distributions of the simulation results.

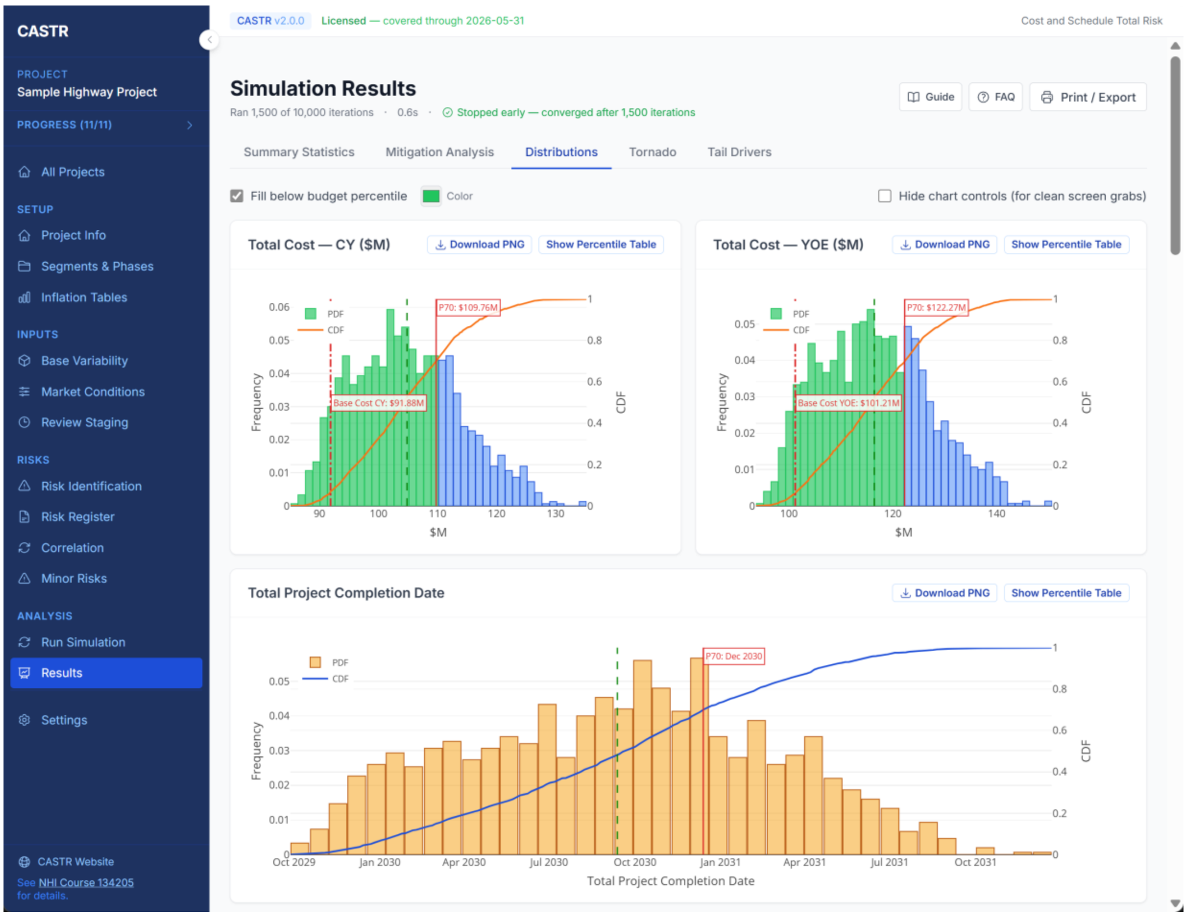

Figure 7.2 - Results Distribution Charts

Three distribution charts are provided:

- **Total Cost (Current Year):** The probability distribution of total project cost in current-year dollars.
- **Total Cost (Year of Expenditure):** The probability distribution of total project cost adjusted for inflation.
- **Total Project Completion Date:** The probability distribution of the overall project completion date.

Each chart displays a histogram showing the probability density function (PDF) with a cumulative distribution function (CDF) overlay curve. The CDF curve allows you to read off the probability of the project cost or completion date falling below any given value.

Chart options include the ability to fill below the budget percentile threshold, show or hide chart controls, download the chart as a PNG image, and display a Percentile Table showing exact values at each tracked percentile.

Reconstructed distributions for SoSP / SoPSP (added in v2.2). When the Budget Percentile Method on the Simulation page is set to Sum of Segment Percentiles (SoSP) or Sum of Phase-Segment Percentiles (SoPSP), the headline budget value is the SUM OF PERCENTILES across components rather than a single percentile of the project total. Under those methods the project-total chart on this tab shows a RECONSTRUCTED distribution that reconciles with the SoSP / SoPSP budget by construction — the reference line lands exactly at the chosen budget percentile of the chart. Per-phase-segment drill-down charts still show their own raw simulated distributions, unchanged. Under PoS the project-total chart is the raw simulated total, as before.

7.3 Tornado Charts

Tornado charts provide a sensitivity analysis showing which risk variables have the greatest influence on the total project cost and schedule. They are essential for identifying where risk mitigation efforts will have the most impact.

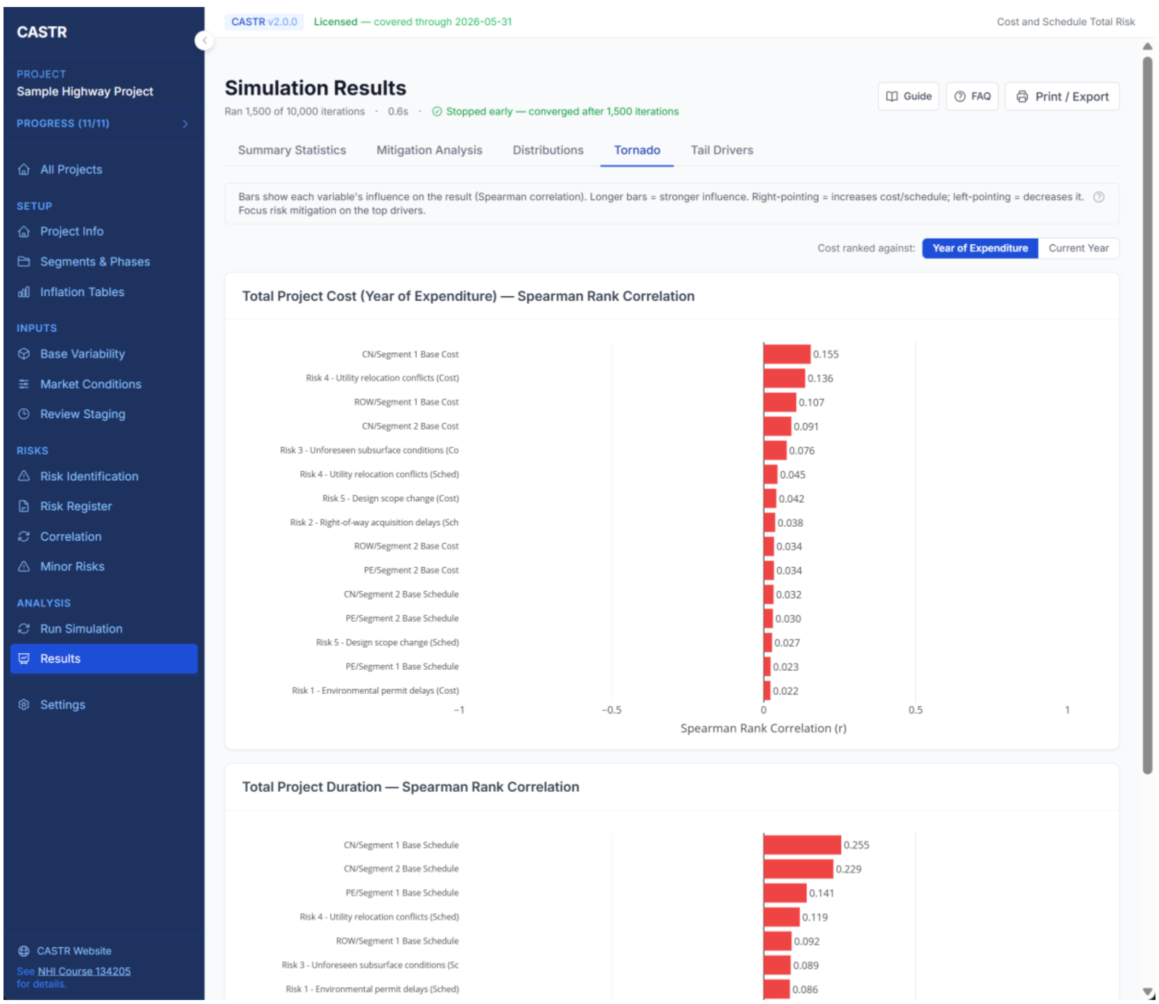

Figure 7.3 - Results Tornado Charts

The charts display horizontal bars representing the Spearman rank correlation coefficient of each variable against the total cost or total duration. Longer bars indicate stronger influence on the outcome.

- Right-pointing bars (positive correlation) indicate that higher values of the variable tend to increase the total cost or duration.
- Left-pointing bars (negative correlation) indicate that higher values of the variable tend to decrease the total cost or duration.

You can toggle between Year of Expenditure and Current Year views. Separate tornado charts are provided for Total Project Cost and Total Project Duration, allowing you to identify the key drivers of cost risk versus schedule risk independently.

Risk Impact mode (added in v2.2). The tornado has a toggle at the top right that switches between two ranking modes. “Spearman” (the default) is the classical sensitivity view — bars show each input's Spearman rank correlation with the simulated outcome and answer the question “which risks drive variability?”. “Risk Impact” ranks risks by their CONTRIBUTION TO MEAN COST in dollars (P(occur) × expected cost impact, signed for threats vs opportunities) and answers the more board-friendly question “which risks add the most expected dollars to the headline estimate?”. Use Spearman to chase tail risk and Risk Impact to chase the central estimate; many CASTR users include both views in their reports.

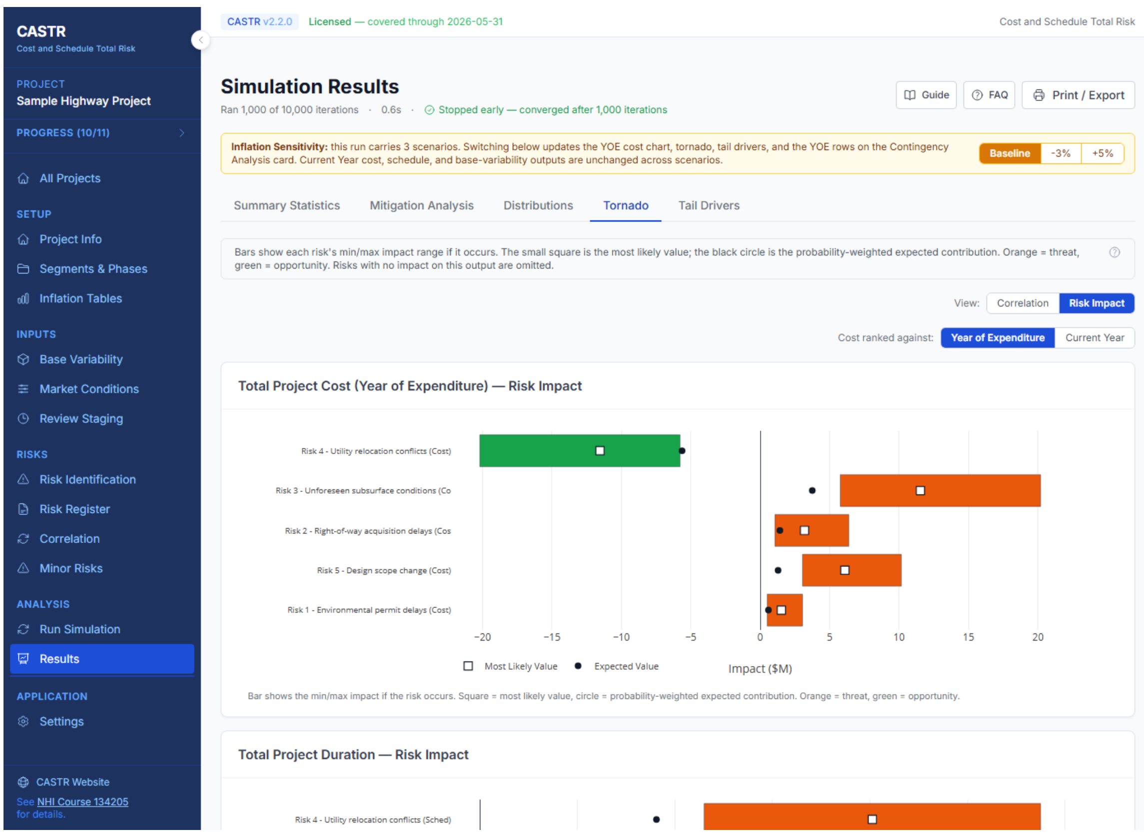

Figure 7.3a — Tornado in Risk Impact mode. Bars are signed expected dollar contributions to mean cost, ranked largest first. Use the toggle in the upper right of the chart to flip between Spearman correlation and Risk Impact.

7.4 Tail Driver Analysis

The Tail Driver Analysis identifies the specific risks that are disproportionately responsible for the extreme outcomes in the simulation. Understanding tail drivers helps project teams focus risk mitigation on the events that contribute most to worst-case (and best-case) scenarios.

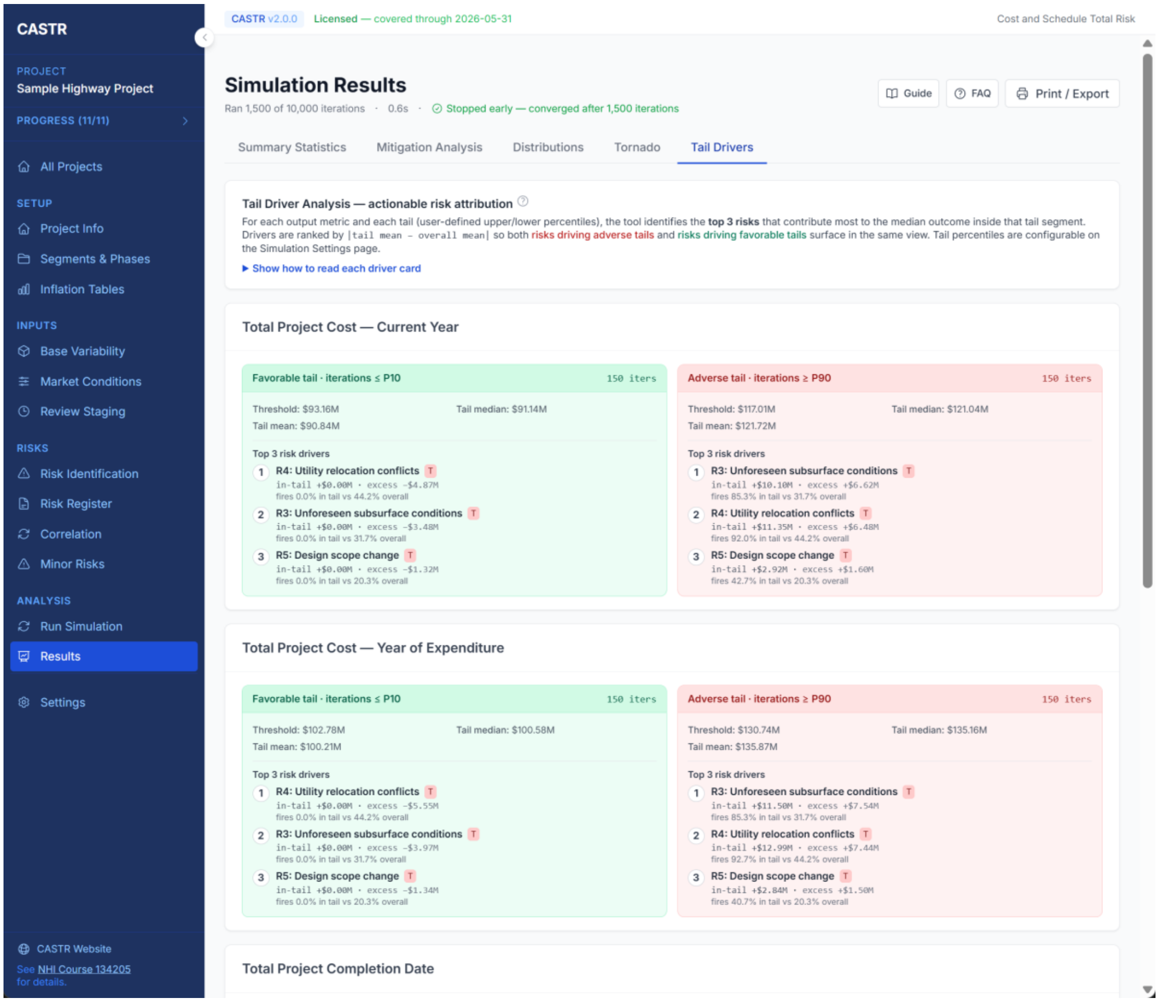

Figure 7.4 - Results Tail Driver Analysis

The analysis identifies the top three risks driving the favorable (low) tail and the top three risks driving the adverse (high) tail. Results are shown for Total Project Cost (both Current Year and Year of Expenditure) and Completion Date.

Each tail driver card displays:

- **Fire Rate in Tail vs. Overall:** How much more (or less) frequently the risk occurs in the tail compared to the overall simulation. A risk that fires 80% of the time in the adverse tail but only 40% overall is clearly a tail driver.
- **Excess Contribution:** The additional cost (in $M) or schedule impact (in months) that the risk contributes in the tail beyond its average contribution.
- **Interpretation:** A plain-language explanation of the risk's role in driving extreme outcomes.

A phase-segment drill-down feature is available to examine which geographic segments are most affected by each tail driver, providing additional granularity for targeted risk response planning.

7.5 Inflation Sensitivity Scenarios on Results (added in v2.2)

When a simulation run included one or more configured Inflation Sensitivity Scenarios (§6.1), the Results page exposes three additions for comparing scenarios apples-to-apples.
Scenario chip selector. A row of chips at the top of every Results tab lists the Baseline plus each enabled scenario by its label (or by the auto-generated delta string when the label is blank). Click a chip to swap the entire view to that scenario. The distribution charts, tornado, tail-driver cards, and pre-vs-post mitigation overlay all repaint per scenario; the selector is sticky, so flipping scenarios does not lose the current tab.
Side-by-side comparison table. The CASTR Contingency Analysis card on the Summary tab includes an Inflation Sensitivity comparison row block that shows Base Cost YOE, Mean Total Project Cost YOE, and the budget-percentile YOE value across every scenario in a single table. This is the executive-defense view — it summarises in one place how the YOE budget shifts when the inflation assumption shifts.
Sensitivity Curve. Below the comparison table is a small Plotly chart that anchors a budget-percentile YOE point at each scenario's delta and interpolates between them with a quadratic Lagrange fit (3 anchors → exact quadratic; 2 → linear; 1 → flat). The domain is auto-extended ~50 % beyond the bracket of your scenarios (capped at ±10 pp). Inside the bracket the curve is the simulation result; outside the bracket it is directional only and the chart caption notes the extrapolation.
Schedule charts (project-completion date, total project duration) are scenario-invariant by design — inflation deltas only affect cost, not schedule — and are deliberately rendered identically for every scenario so that flipping the chip selector confirms the schedule basis is unchanged.

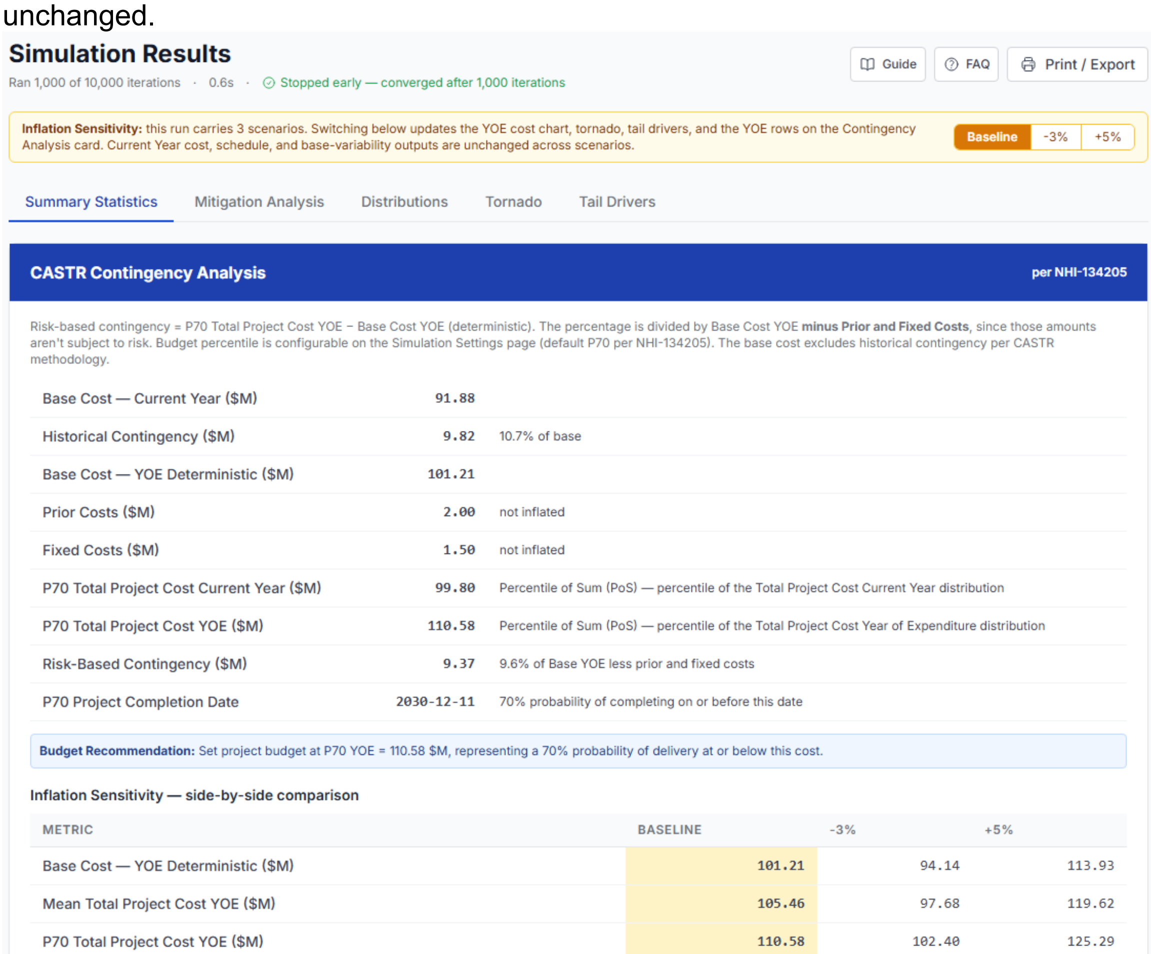

Base Cost — Current Year ($M)	91.88	
Historical Contingency ($M)	9.82	10.7% of base
Base Cost — YOE Deterministic ($M)	101.21	
Prior Costs ($M)	2.00	not inflated
Fixed Costs ($M)	1.50	not inflated
P70 Total Project Cost Current Year ($M)	99.80	Percentile of Sum (PoS) — percentile of the Total Project Cost Current Year distribution
P70 Total Project Cost YOE ($M)	110.58	Percentile of Sum (PoS) — percentile of the Total Project Cost Year of Expenditure distribution
Risk-Based Contingency ($M)	9.37	9.6% of Base YOE less prior and fixed costs
P70 Project Completion Date	2030-12-11	70% probability of completing on or before this date

Budget Recommendation: Set project budget at P70 YOE = 110.58 $M, representing a 70% probability of delivery at or below this cost.

Inflation Sensitivity — side-by-side comparison

METRIC	BASELINE	-3%	+5%
Base Cost — YOE Deterministic ($M)	101.21	94.14	113.93
Mean Total Project Cost YOE ($M)	105.46	97.68	119.62
P70 Total Project Cost YOE ($M)	110.58	102.40	125.29

Figure 7.5 — Results Summary with the scenario chip selector (top), the CASTR Contingency Analysis card, and the side-by-side Inflation Sensitivity comparison table. Switching the chip selector repaints every other tab on this page accordingly.

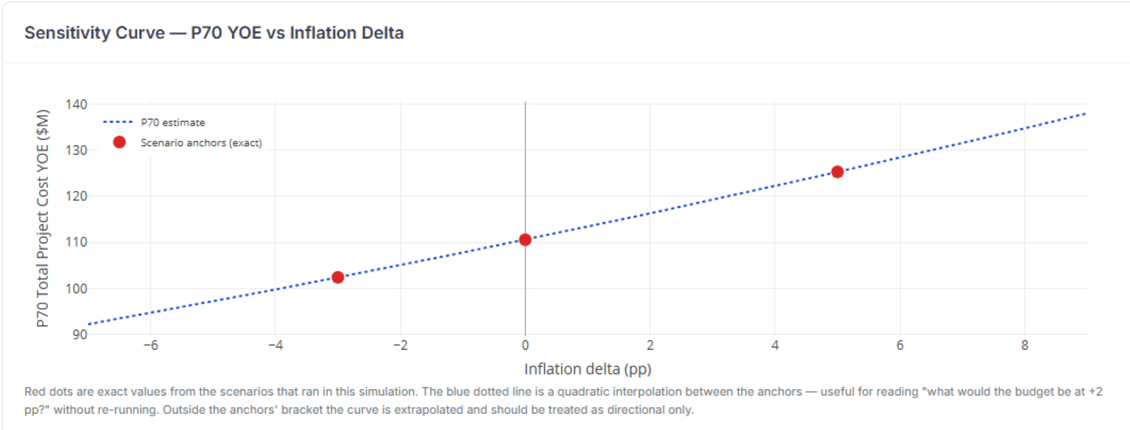

Figure 7.5a — Sensitivity Curve. Red dots are exact simulated values at each scenario delta; the dotted line is the quadratic interpolation between them. Outside the bracket the curve is extrapolated and should be treated as directional.

7.6 Word Export (added in v2.2)

The Print / Export button on the Results header offers a one-click Export to Word option that produces a polished .docx report. The exported document includes the Project Information block, summary statistics, the CASTR Contingency table, the Mitigation ROI tables, and embedded chart images (CDF, histogram, schedule histograms, tornado, completion date, total project duration). Charts are rendered at presentation-friendly font sizes (axis ticks ≈22 pt, axis titles ≈26 pt, annotations ≈22 pt, legend ≈20 pt) so the labels stay legible when the report is shared, projected, or printed.

When the run includes Inflation Sensitivity Scenarios (§6.1) the export modal shows an “Include all inflation scenarios” checkbox. With it ticked, the Distributions and Tornado sections are duplicated once per scenario with amber-bordered scenario headings; CY, schedule, and completion-date sections (which are scenario-invariant) print once at the end. With the box unticked only the currently-selected scenario is exported. The browser-print modal mirrors the same per-scenario behaviour for paper or PDF audit trails.

Chapter 8: Data Import and Export

8.1 Risk Import from Excel

New in v1.8: The Excel import template now includes 13 additional mitigation columns (Risk Response, Mitigation Cost, Owner, Due Date, P(RESID), and the residual cost/schedule distribution parameters). Leave them blank on legacy templates — existing workflows continue to work unchanged. New in v1.9: a dedicated Checkpoints dialog lets you save and restore named snapshots of the risks and register rows at any time (see §8.4). Every destructive operation (Import, Clear All, Restore, Bulk Delete) still takes an automatic safety snapshot, so you can always roll back.

CASTR supports importing risk data from a standardized Excel template. This feature is particularly useful when risks are identified during workshop sessions and recorded in spreadsheets, or when migrating data from other risk management tools.

To import risks from Excel:

1. On the Risk Identification page, click "Import from Template..." to open the import dialog.
2. Download the pre-configured Excel template using the provided link. The template includes data validation dropdowns populated with your project's categories and phase-segments.
3. Fill in the risk data in the template. Each row represents one risk event.
4. Upload the completed template back into CASTR using the import dialog.
5. Review the preview showing valid rows, warnings, and errors. Resolve any errors before proceeding.
6. Choose Append (add imported risks to existing data) or Replace (clear all existing risks and import fresh). Click Import to complete the process.

A snapshot of your existing risk data is automatically created before any destructive import operation (Replace mode), allowing you to undo the import if needed.

8.2 Project Backup and Restore

The Backup and Restore feature provides a reliable way to preserve project data, transfer projects between installations, and recover from data issues.

Exporting Projects: From the home page, click "Backup / Restore" and select the Export tab. Check the boxes next to the projects you want to export, then click the download button. All selected projects are packaged into a single .zip file that contains the complete project data, including all settings, risk entries, correlations, inflation tables, market conditions, and simulation configurations.

Importing Projects: Select the Import tab and upload a previously exported .zip file. The application displays a preview of the projects contained in the file. If any project names match existing projects, you can choose a resolution strategy for each conflict: skip (keep the existing project unchanged), rename (import with a modified name), or overwrite (replace the existing project with the imported version). Click Commit to finalize the import.

Full project data is preserved through the export/import cycle, ensuring no loss of information when transferring projects between machines or creating backups.

8.3 CSV/Excel Export from Risk Register

The Risk Register page provides direct export functionality for generating reports in common spreadsheet formats. Use the CSV button to export the risk register data as a comma-separated values file, or the Excel button to export as a formatted .xlsx file.

The exported file includes all columns currently visible in the selected view preset (Brief, Standard, Full, or Custom). This makes it easy to generate tailored reports for different audiences by selecting the appropriate view preset before exporting.

8.4 Checkpoints

Overview

Checkpoints (added in v1.9) are named, point-in-time snapshots of a project’s Risk Identification and Risk Register data. Their primary purpose is to support iterative workshops: create a checkpoint before an SME session (e.g. “Before SME Workshop – 2026-04-17”), iterate freely during the meeting, and restore the checkpoint if the group wants to walk back a block of changes. Every Restore operation itself takes an automatic safety snapshot named “Before restore of #N”, so restores are always reversible.

Creating, Restoring, and Deleting

A Checkpoints button is available at the bottom of the Risk Identification tab and below the Risk Register table. Clicking it opens the Checkpoints dialog. Type a descriptive label in the input field (or leave blank to auto-stamp with the current date and time) and click Save Checkpoint. Existing checkpoints are listed below with their label, creation timestamp, and the number of risks captured; the Restore link rewrites the risks and register rows back to that point, and Delete removes the checkpoint permanently. Restoring triggers a confirmation prompt and notes that a safety snapshot is taken first.

Automatic Snapshots and the 25-Checkpoint Cap

Destructive operations automatically take their own checkpoints before running: Replace-mode import, Clear All Risks, Restore of a prior checkpoint, and the new Bulk Delete (multi-row delete from the SEL column on Risk Identification). These automatic snapshots appear in the same checkpoint list with system-generated labels (“Before Clear All …”, “Before bulk delete …”, etc.) so you can roll any of them back from the same dialog. A project is capped at 25 user-created checkpoints to keep the list manageable and storage bounded. When the cap is reached the Save Checkpoint button disables with an amber warning; delete an older checkpoint to free a slot. Automatic safety snapshots are exempt from the cap so destructive operations always retain their undo path.

Scope and Limitations

A checkpoint captures the risk_events and risk_register_rows tables for one project, including all MIT flags, status, locked states, and row numbering. It does not capture Base Variability, Market Conditions, Inflation Tables, Correlations, or simulation runs — those are treated as setup data and iterated separately. Checkpoints are stored locally in the castr.db file and are intentionally excluded from Backup/Restore exports (Section 8.2), so exported project files are small and transportable; if you want to share a snapshot with another user, restore it first and then export the project.

Chapter 9: Starter Risk Library

9.1 Overview

The Starter Risk Library is an app-global list of predefined risk categories and risks that every project on this CASTR install shares. It ships pre-populated with a default set of highway construction risk categories (Third-Party, Construction, Design, Environmental, etc.) and ~100 common risks you can quickly drop into a new project instead of typing every row from scratch.

The library is fully editable through a dedicated admin page. Any edits you make apply to every project — they appear in the "Starter Library…" picker on the Risk Identification page the next time you use it. Rows already inserted into a project are independent and are NOT affected by later edits to the library.

9.2 Editing the Library

Click "Starter Library" in the left sidebar (below "All Projects"). The page has two tables stacked vertically:

Categories — 1–4 letter Cat ID plus a free-text description. Sortable by Cat ID or Description (click the column header to toggle asc / desc / no-sort).

Risks — Cat ID (dropdown), Risk Event Name, and Risk Description. Sortable by Cat ID. Includes a chip bar above the table for filtering risks by one or more categories (counts update live).

All fields edit inline and save on blur. A blank field is backshaded light yellow to draw the user's attention — this is purely a reminder, not a blocker.

Click "Add Category" or "Add Risk" at the bottom of each table to create a new row. New Add-Risk rows default their Cat ID to whichever category filter is currently active, so the row is immediately visible after creation. The trash icon on each row deletes that entry; deleting a category does not delete the risks that reference it — they simply show as "orphan" until you re-categorize them or delete them.

Excel Import (added in v2.2). The “Import from Excel” button at the top of the Starter Library page opens a three-step wizard. Step 1 — download the template, an .xlsx workbook with two sheets (Categories: Cat ID + Description; Risks: Cat ID + Risk Name + Description) pre-filled with the current library so you can edit in place. Step 2 — pick your filled file; the server validates each row and shows a preview listing how many categories and risks will be added or updated, plus any rows that could not be parsed. Step 3 — click Commit (append) to apply. The import is additive and idempotent: existing categories with matching IDs are updated in place, new ones are appended, and starter risks are matched on (Cat ID, Risk Name) so re-importing the same file is safe.

Restore Defaults (added in v2.2). The “Restore Defaults” button next to Import rebuilds the library from the factory NHI / NCHRP starter catalog that ships with CASTR. Use it to recover from a botched bulk edit or import. Restore does NOT touch projects that already imported starters — those stay frozen at their imported values, consistent with the snapshot semantics described in §9.1. Warning: Restore Defaults REPLACES the entire current Starter Library — every category and every starter risk currently saved will be overwritten. The action prompts for confirmation before proceeding.

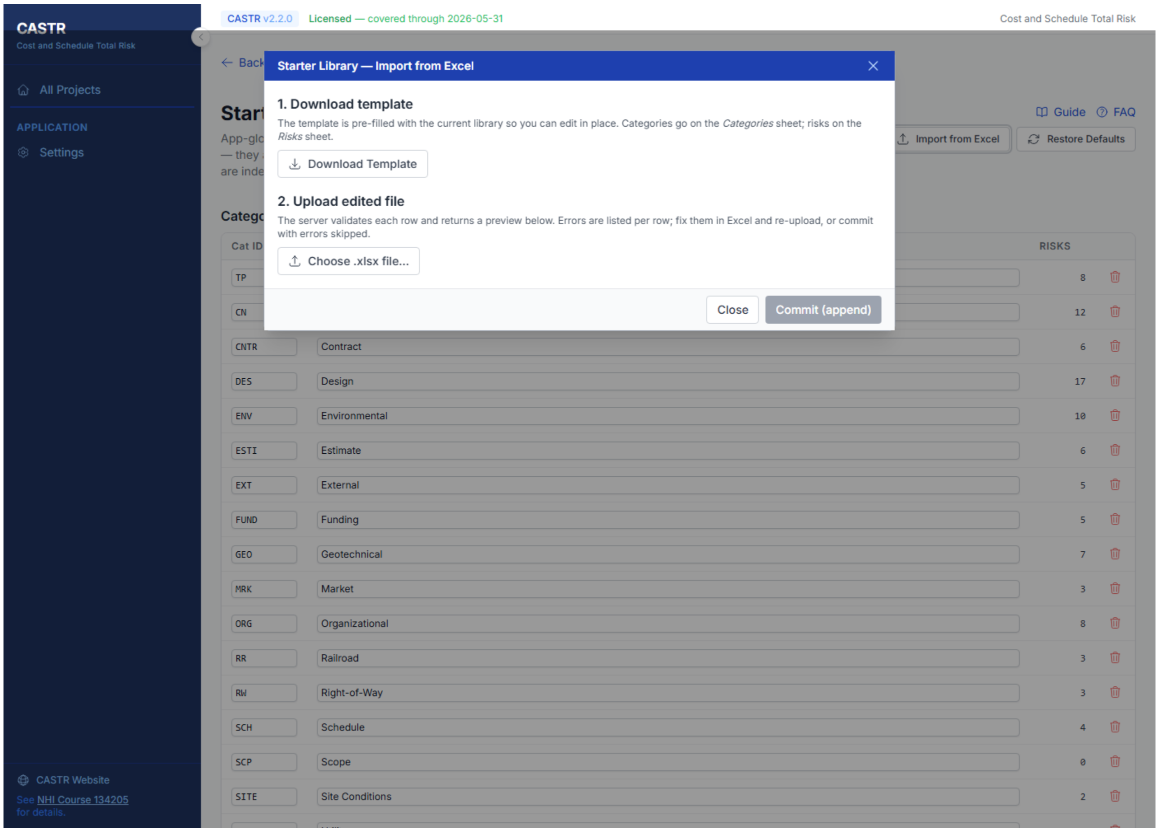

Figure 9.1 — Starter Library Import-from-Excel wizard. Step 1 downloads a pre-filled template; step 2 uploads the edited workbook and shows a preview; step 3 commits the changes.

9.3 Relationship to Project Risk Categories

The Starter Library's Categories table is separate from the per-project Risk Categories table shown on the Segments & Phases page. The two are independent, but they connect in one direction: when you pick a starter risk from the Risk Identification page, if its Cat ID is not already present on that project, the category is copied into the project's Risk Categories table automatically. After the one-time copy, the two are fully independent — editing a starter category later does not change any project's local category, and vice versa.

9.4 Using the Starter Library Picker

On the Risk Identification page, click the "Starter Library…" button next to "Add Risk" and "Import from Template…". A modal opens with two columns:

Left column — every starter category with a live count of risks in it. Click a category to toggle it on/off; select multiple categories with OR logic, or click "All categories" to clear the selection. A green checkmark on a category row means that category already exists on the project.

Right column — every starter risk, filtered by the selected categories and by the search box at the top. Each row has a checkbox; select any number of risks, then click "Add N to Risk Table".

Selected risks are appended as new rows on the Risk Identification table. Any categories that were referenced but not yet on the project are created first, automatically. The alert confirms both counts ("Added 4 starter risks (2 new categories added)"). Inserted rows are regular Risk Events — fully editable, renumbered, synced, snapshotted, and backed-up exactly like manually created ones.

Chapter 10: Risk Mitigation Analysis

10.1 Overview

CASTR v1.8 introduced a full mitigation workflow that lets you capture the planned response to each risk, the cost of the mitigation action, the residual probability and impact after the action is applied, and then run the Monte Carlo simulation in either pre-mitigation, post-mitigation, or compare-both mode to quantify the return on the mitigation program. Version 1.9 rounds this out with a workshop-style Residual Distribution Editor — it opens with a green backshade and an amber “MIT” badge in the header to clearly distinguish it from the original distribution editor, shows every original parameter as ghost text under its corresponding residual input so the subject-matter expert can calibrate against the baseline, renders the preview chart automatically (original curve in red, residual curve in green) without having to click Preview, and is draggable by its header so you can move it aside to reference the underlying Risk Details dialog. Residual P(RESID) is now required whenever Risk Response is set to Mitigate: attempting to save a blank value turns the field light red and blocks the save with a specific error message. The Sample Highway Project now ships with two fully-mitigated risks (Environmental permit delays and Utility relocation conflicts) with MIT toggled on, so a new installation can be simulated end-to-end in Pre, Post, and Compare-Both modes out of the box; existing installations are auto-patched once on the next startup.

Define every mitigation field once on the Risk Identification page (Risk Details panel). Sync carries them forward into the Risk Register, where a dedicated MIT checkbox controls whether each risk's residual parameters are used by the simulator. The Results page shows a Mitigation ROI Summary whenever any risk has MIT active, and overlays pre vs post curves when the run mode is Compare Both.

10.2 Defining Mitigation on a Risk

Tip (added in v1.10): the Risk Identification Guide panel now includes a dedicated "Setting Mitigation - Step by Step" section that walks through the four-step workflow (open Mit/View dialog -> choose Mitigate and fill Mitigation Cost + Residual P(RESID) -> Edit Residual Cost/Schedule Distribution -> flip MIT on the Risk Register). It explicitly warns that stopping before step 4 captures the plan but does NOT change simulation output. A companion "Mitigation - Visual Cues to Watch For" section describes the green-shield icon, yellow-backshaded residual inputs, and the authoritative role of the MIT checkbox. Open the Guide panel on Risk Identification to see them.

On the Risk Identification table, click the View (or Mit) button in the Details column for the risk you want to mitigate. This opens the Risk Details panel.

Risk Response — the Risk Response field is a dropdown with four options:

Mitigate — an action will be taken to reduce the probability or impact of the risk. Unlocks the Mitigation Details section below.

Avoid — the risk is eliminated by changing the project approach.

Transfer — the risk is assigned to a third party (insurance, subcontractor, etc.).

Accept — the risk is acknowledged but no action is planned. This is the default for legacy risks and new risks.

When Risk Response = Mitigate, a Mitigation Details section appears with the following fields:

Mitigation Cost ($M) — the certain cost (100% probable) of performing the mitigation action. Added to the base estimate when the simulation runs in post-mitigation mode.

Residual P(RESID) — the probability the risk occurs AFTER mitigation, between 0 and 1.

Mitigation Owner — the person or role responsible for executing the mitigation action.

Mitigation Due Date — when the mitigation action must be complete. The Risk Register highlights past-due dates when MIT is active.

Edit Residual Cost Distribution — opens the distribution editor against a dedicated residual cost distribution. The first time you open it, residual parameters are pre-populated from the original cost distribution so you start from a sensible baseline and only adjust what changes.

Edit Residual Schedule Distribution — same as above but for the schedule impact.

The residual distribution editor is the same component used for the original distribution — you can switch between Cost and Schedule tabs, preview the curve, and adjust the distribution type (Triangular, BetaPERT, Uniform, Normal). The Additive-with-other-schedule-risks checkbox applies in both pre- and post-mitigation modes; it is a property of the risk, not of the distribution.

Visual indicator — once Risk Response = Mitigate is saved, the Details-column button for that risk turns green and changes label from "View" to "Mit", with a shield icon. This makes mitigated risks visible at a glance when scanning the Risk Identification table.

Residual P(RESID) reference cues — when this field is blank the editor shows the risk’s current P(occur) as placeholder text, backshades the input yellow, and repeats the value in a hint line below the field (e.g. “Current P(occur) is 0.35”). All three cues clear as soon as you enter a residual value, but the hint line stays visible so you can always compare the pre- and post-mitigation probabilities side by side without closing the dialog.

Draggable title bar — click and hold the Risk Details title bar to drag the dialog anywhere on screen. Useful when you want to see the underlying Risk Identification row (for example, to cross-check the original cost impact while filling in the residual distribution). The dialog re-centers automatically each time you re-open it.

Blank-field highlights on the Mitigation Details section — while the Mitigation Cost field is empty the input is backshaded light yellow, and the “Edit Residual Cost Distribution” and “Edit Residual Schedule Distribution” buttons are backshaded yellow until the corresponding residual distribution has been configured. The cues match the P(RESID) reference cue above: they signal “this is something that still needs your attention” and clear automatically as soon as a value is saved. Hover tooltips explain why each cue is on.

Pre vs Post preview overlay — when the Residual Distribution Editor is open and you click “Preview Distribution”, the chart renders four traces: the original (pre-mitigation) PDF and CDF in red, and the residual (post-mitigation) PDF and CDF in green. The two colors reuse the Pre/Post convention from the Results overlay so the subject-matter expert can see both curves on the same axes and pick residual parameters by comparison. Changing the residual parameters on the left and clicking Preview again refreshes both curves side-by-side. The original curve is read off the unchanged risk record — editing the residual never shifts the baseline.

10.3 Filtering by Mitigation Status

The Risk Identification filter bar includes a three-way segmented control: All / Mit / Non-Mit. Each segment shows a live count of risks in that bucket within the current category filter. Use it to quickly isolate the mitigated risks for review, or to see what's left that hasn't been assigned a mitigation plan.

10.4 Sync to Risk Register and the MIT Column

When you click "Sync to Risk Register", every Risk Identification row becomes a Risk Register row with a matching Reg # (reg # always equals row #). All mitigation fields are carried forward via the sync.

The Risk Register gains an MIT checkbox column positioned after Sched Dist. and before Status. This column, along with the other mitigation columns (Response, P(RESID), Mitigation Cost, Mit Owner, Mit Due Date, Residual Cost Dist., Residual Sched Dist.), is hidden by default and only appears when at least one risk on the project has Risk Response = Mitigate.

MIT is the master switch. It tells the Monte Carlo simulator, for this specific risk, to use the residual (post-mitigation) parameters instead of the original ones when the run mode is Post-Mitigation or Compare Both. A risk can be planned for mitigation (Risk Response = Mitigate, residual params filled in) without MIT being checked — the MIT checkbox is the separate commitment step that says "this mitigation is funded and active, use it in simulation".

Editable only when Risk Response = Mitigate for that risk; dimmed and non-interactive otherwise.

Protected from sync overwrite — once the user sets MIT, re-syncing never clears it. On the first sync of a row, MIT is pre-checked if Risk Response = Mitigate, so the common case needs no extra clicks.

Past-due highlight — when MIT is active and the Mitigation Due Date is before today, the date cell is highlighted in red.

Orphan preservation — if a risk is deleted from Risk Identification but its register row has MIT active or Status ≠ Active, the register row is preserved with an orphan flag instead of being silently deleted. An amber banner above the table summarizes the orphan count.

10.5 Run Mode

The Simulation page has a Run Mode card with three radio options:

Pre-Mitigation (Inherent Risk) — the default. Uses P(occur) and the original cost/schedule distributions for every risk. MIT is ignored; results are byte-for-byte identical to historical CASTR behavior.

Post-Mitigation (Residual Risk) — for each risk with MIT active: substitutes P(RESID) for P(occur), substitutes the residual cost/sched distributions for the originals, and adds the sum of Mitigation Cost values to the project base estimate as a certain (non-random) cost. MIT-unchecked risks still use their original parameters.

Compare Both — runs the simulation twice, once in Pre-Mitigation and once in Post-Mitigation mode, and passes both complete result sets to the Results page. Roughly 2× wall time.

When MIT is unchecked on every risk, Post-Mitigation and Compare Both produce the same numbers as Pre-Mitigation (the parameter-swap path never triggers). This is harmless but wastes compute time, so switch to Post-Mitigation or Compare Both only after you have ticked MIT on the risks whose mitigation is funded.

10.6 Mitigation ROI Summary and Charts

The Mitigation Analysis tab appears on the Results page between Summary Statistics and Distributions whenever the run produced mitigation data — either a ROI summary (because at least one risk on the register has MIT active) or a post-mitigation result set (because Run Mode was Compare Both). For projects that are not using mitigation, the tab does not appear at all and the Results page shows the same four tabs as before.

The tab contains two sections: a Mitigation ROI Summary table that is always shown when the tab is visible, and a Pre vs Post Comparison block that appears only for Compare-Both runs.

Mitigation ROI Summary — a per-risk table with these columns:

Reg # — register row number.

Risk Event — the risk name.

Mitigation Cost — the certain cost of the mitigation action.

Pre-MIT Expected Value — P(occur) × Most-Likely cost impact (from the original cost distribution).

Post-MIT Expected Value — P(RESID) × Most-Likely residual cost impact (from the residual cost distribution).

Net Benefit — Pre-MIT EV − Post-MIT EV − Mitigation Cost. Positive means the mitigation is worth its cost; negative means it costs more than the expected risk exposure it removes.

ROI — Net Benefit ÷ Mitigation Cost, displayed as a ratio (e.g. 12.4x). When Mitigation Cost is zero, ROI is shown as "—" to avoid implying an infinite return.

A totals row at the bottom aggregates Mitigation Cost, Pre-MIT EV, Post-MIT EV, and Net Benefit, with an overall ROI = Total Net Benefit / Total Mitigation Cost.

Pre vs Post Comparison charts — shown only when Run Mode = Compare Both. The charts render in a 2×2 grid (collapsing to a single column on narrow windows and in print):

• Top-left: Total Project Cost, Current Year ($M) — pre-inflation total.

• Top-right: Total Project Cost, Year of Expenditure ($M) — inflated.

• Bottom-left: Total Project Duration (months).

• Bottom-right: Total Project Completion Date (requires base-uncertainty rows on the project; otherwise the quadrant shows an explanatory empty state).

Red is Pre-Mitigation, green is Post-Mitigation, and a legend in the top-left of each chart identifies the two curves. There is no Cost-axis dropdown: all four charts are visible simultaneously so the user can see cost and schedule effects side by side.

Budget-percentile markers on every comparison chart — when the Simulation page has "Show on chart" enabled for the budget percentile AND the budget method is Percentile of Sum (PoS), each of the four overlays draws two vertical dashed lines (red at the Pre value, green at the Post value at that P-level) with compact labels showing the value, and a caption below the chart reports the difference at the budget level.

Label and caption formats adapt to each chart's unit: cost charts use dollars-millions (e.g. "Pre P70 $101.40M", "Δ @ P70: −$0.17M"); the duration chart uses months ("Pre P70 119.64 mo", "Δ @ P70: −2.65 mo"); the completion-date chart uses ISO dates with a day-count delta ("Pre P70 2030-07-01", "Δ @ P70: −52 days"). A minus sign in the delta indicates the Post value is lower/earlier than the Pre value — i.e. the mitigation improved the outcome at the budget level.

Markers are not rendered when show_budget_percentile is off, when the method is SoSP or SoPSP (those methods are sums of percentiles and do not correspond to a single point on the total distribution), or when the data is unavailable. When enabled, the markers appear on all four comparison charts.

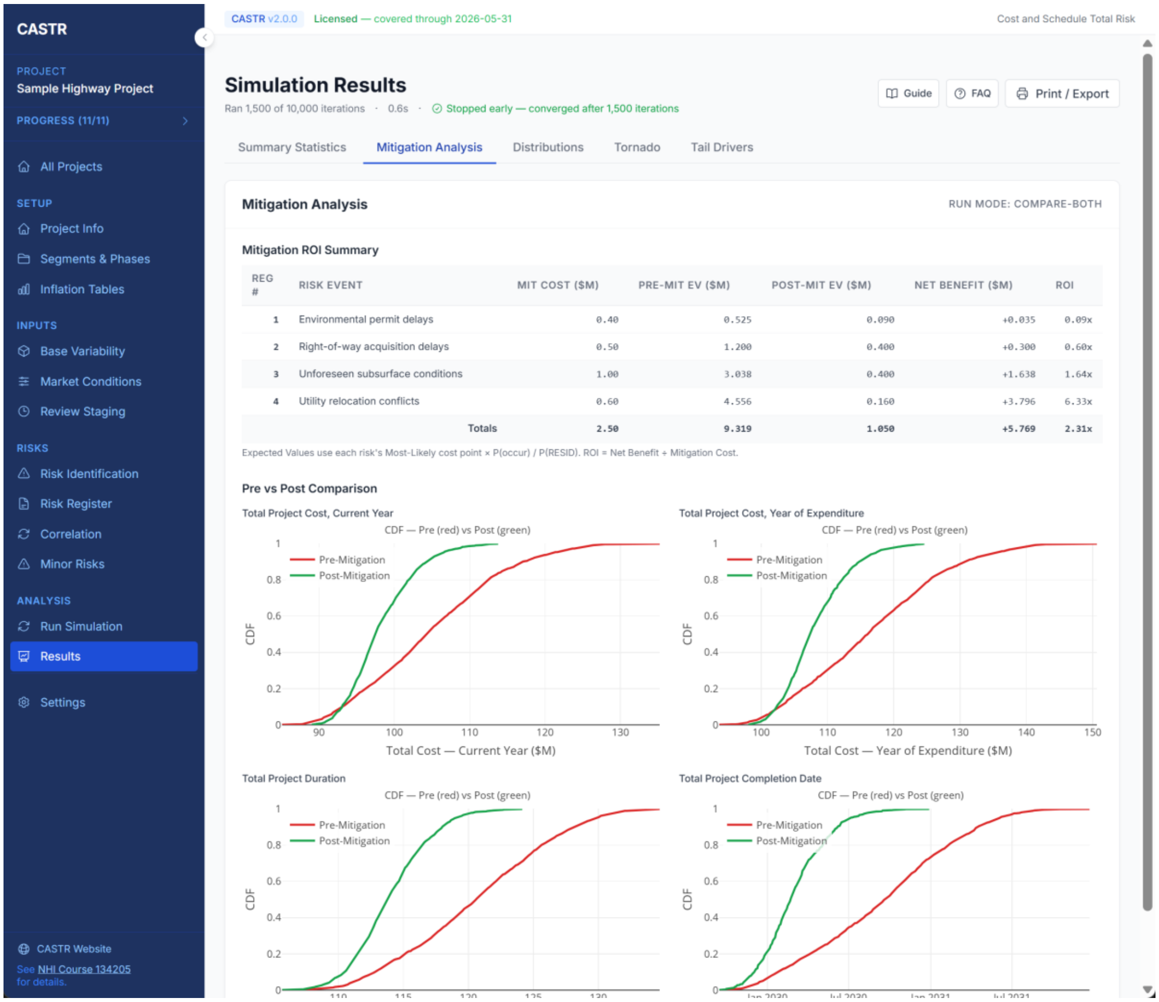

Figure 10.1 - Mitigation Analysis: pre- vs post-mitigation comparison of the project cost distribution, with mitigation cost and ROI per risk summarized below.

Mitigation ROI — by Inflation Scenario (added in v2.2). When the run includes Inflation Sensitivity Scenarios (§6.1) AND the Risk Register has at least one active mitigation, a second table titled “Mitigation ROI — by Inflation Scenario (YOE)” appears below the headline CY ROI table. It reports Pre-MIT EV, Post-MIT EV, Net Reduction, and ROI per risk under each scenario's per-risk inflation factor in YOE dollars. Mitigation Cost is NOT scaled by inflation — mitigation spend is committed in current-year dollars, and the CY-cost vs YOE-EV ratio is the defensible apples-to-apples ROI reading. The CY ROI table is unchanged so the headline ROI number remains stable across scenarios; the YOE sub-table is for showing executives how the size of the avoided-cost pie shifts under different inflation assumptions.

10.7 Quick Reference

Typical end-to-end workflow for mitigation:

Step 1 — On Risk Identification, open a risk's Details panel and set Risk Response to Mitigate. Fill in Mitigation Cost, Owner, Due Date, P(RESID), and click the two Edit Residual Distribution buttons to tune the post-mit curves. The Details button for this row turns green.

Step 2 — Click Sync to Risk Register. On first sync only, MIT is auto-checked on that register row because Risk Response = Mitigate.

Step 3 — Optional: on the Risk Register, uncheck MIT for any mitigations that are planned but not yet funded.

Step 4 — On Simulation, pick Run Mode = Post-Mitigation or Compare Both. Run.

Step 5 — On Results, review the Mitigation ROI Summary table and (for Compare Both) the overlay CDF curves.

Chapter 11: Global Scoring Tables (Settings)

11.1 Overview

Four tables drive the starter distributions produced by the Seed action and the Contingency Envelope check. All four are editable on the Settings → Scoring Tables page (added in v1.9). The defaults are drawn from the NHI Course 134205 / NCHRP Report 6-17 methodology and are appropriate for typical highway construction projects in the United States. Agencies with different standards or specialized project types may adjust them to fit.

All four tables are global — changes affect every project on this CASTR install. Edits take effect immediately for subsequent Seed operations but do NOT retroactively change risk distributions that have already been seeded or entered. If you need to re-apply the new table values to an existing risk, open that risk and click Seed again (confirming the overwrite in the dialog).

Access: click Settings in the left sidebar, then click the Scoring Tables card. Each table displays a badge showing Default values (no user edits) or Custom (edited). Save persists the current table to the local database. Restore defaults removes any user override and reverts the table to the built-in values.

11.2 Cost Scoring Table

The Cost scoring table maps a cost_score of 1–5 to low / mode / high fractions of the risk's phase-segment Base Cost. When a risk with cost_score = k is seeded, the starter distribution's parameters are:

p1 (low) = Base Cost × cost_bands[k].low × calibration_factor

p2 (mode) = Base Cost × cost_bands[k].mode × calibration_factor

p3 (high) = Base Cost × cost_bands[k].high × calibration_factor

Example: for a phase-segment with a Base Cost of $10 M and a risk with cost_score = 3, the default bands (0.030 / 0.060 / 0.100) produce a raw starter triangle of $0.30 M / $0.60 M / $1.00 M. If the phase-segment is average size for the project, the calibration factor is 1.00 and those same values are written. If the phase-segment is materially smaller than the project average, the factor is pinned at 0.75 and the written values become $0.225 M / $0.45 M / $0.75 M.

Editing tips:

The table requires exactly five rows (scores 1–5). The Save button is disabled if any row is out of order (low > mode or mode > high).

An adjacent read-only column shows each row as a percent, so you can reason about the bands in the familiar 1 % / 6 % / 35 % shape while still editing in decimal form.

Use the Restore defaults button to revert to the NHI 134205 values at any time. A small confirmation dialog protects against accidental clicks when edits are unsaved.

Warning: raising the Cost bands will push every Quantify-flagged risk's implied contingency higher, which can push the Contingency Envelope panel from Within into Approaching or

Exceeded. Consider whether the Design Completion & Contingency Envelope bands (§11.5) also need adjustment to match your agency's standards.

Scaling Factor (added in v1.10): a new control above the table lets you scale every Low / Mode / High cell by a percentage in a single operation. Enter a value other than 100 (e.g., 120 for a 20% uplift or 80 for a 20% tightening) and click Apply - every numeric cell is multiplied by the factor, the input resets to 100, and the Save button becomes enabled so you can persist or discard. This avoids per-cell editing when the whole scoring scheme needs a proportional adjustment to match your agency's historical bid data.

11.3 Schedule Scoring Table

The Schedule scoring table maps a schedule_score of 1–5 to low / mode / high extra weeks. Unlike cost, schedule impacts are absolute durations — they do not scale with a base value. When a risk with schedule_score = k is seeded, the starter distribution's parameters are (sched_bands[k].weeks ÷ 4.3482) months, optionally scaled by the same phase-segment size calibration factor used for cost.

Example: for a risk with schedule_score = 4, the default bands (8 / 16 / 24 weeks) convert to 1.84 / 3.68 / 5.52 months. On a longer-than-average phase-segment where the calibration factor is pinned at 1.25, the written values become 2.30 / 4.60 / 6.90 months.

Editing tips:

Values are entered in weeks; the read-only sidebar column shows the equivalent months for sanity checking.

Schedule bands have a very wide legitimate range (a single-week rework vs a year-long permit delay), so there is no upper bound on the values — but Save blocks any row where low > mode or mode > high.

Warning: editing the Schedule scoring table has no effect on the Contingency Envelope panel, which is cost-only. Schedule impact shows up in the simulation's schedule distribution and tornado chart, not in the envelope check.

Scaling Factor (added in v1.10): the same Scaling Factor control appears above the Schedule Scoring Table. It multiplies every Low / Mode / High (in weeks) by the chosen percentage, useful when your agency's typical schedule risk magnitudes differ systematically from the defaults.

11.4 Probability Bands (added in v2.2)

The Probability Bands table maps a Triage probability_score (1–5) to a P(occur) decimal value used by Seed Distributions when the risk's P(occur) cell is empty. The factory defaults follow the AASHTO 1=Very Low … 5=Very High convention: 1=0.10, 2=0.30, 3=0.50, 4=0.70, 5=0.90.

How it is used. On Risk Identification, when a workshop facilitator scores a risk's Prob in the Triage view and the row's P(occur) is still blank, Seed Distributions writes the band value into P(occur) and tags the cell with a yellow shield badge as a “seeded — please confirm” nudge (see §5.1.4). User-typed P(occur) values are never overwritten by Seed; the seeding only fills blank cells. The Triage Prob column header has a “?” help button that opens a panel showing the live mapping, and the 1–5 score chips themselves show the mapped P(occur) on hover so a facilitator can preview what each click will produce.

Editing tips. Each value must be between 0 and 1, and each score must appear exactly once. Mappings should be monotonically increasing across scores (a higher Prob score should imply at least as much probability as a lower score) — the validator does not enforce monotonicity but workshop participants will be confused if the relationship is reversed. Edits are global and not retroactive: they only affect future Seed operations and the live tooltip / help panel.

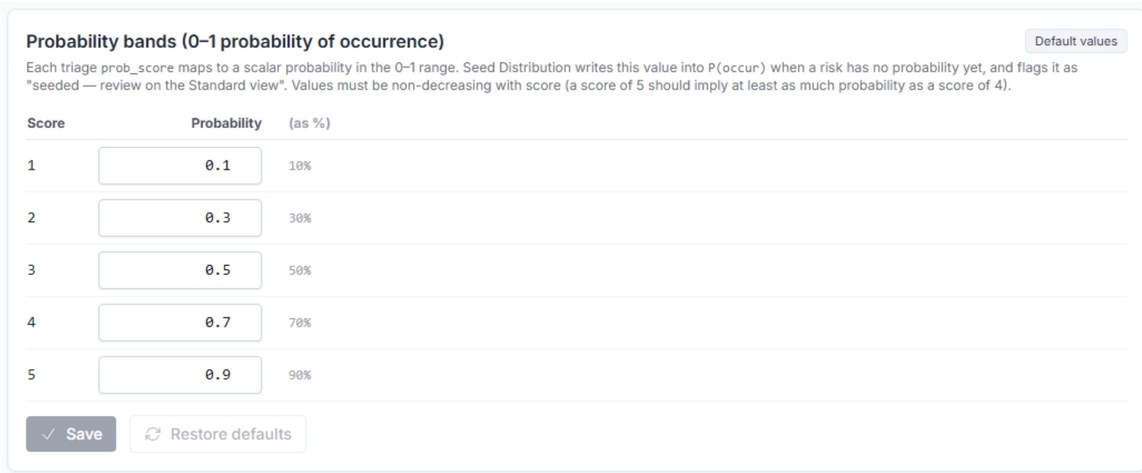

Figure 11.1 — Probability Bands editor. Each Triage probability score (1–5) maps to a scalar probability in the 0–1 range. Edits are global; not retroactive.

11.5 Design Completion & Contingency Envelope Bands

The Design Completion table defines the envelope against which the project's implied whole-project contingency is compared. Each row defines (completion_low, completion_high, contingency_low_pct, contingency_high_pct, label). The row with a completion_low ≤ design % ≤ completion_high interval supplies the envelope range; boundary matches resolve to the more mature band (e.g. exactly 25 % completion lands in the 25–35 % row, not the 10–25 % row).

Default bands (NHI 134205):

1–10 % (Concept): 30 %–75 % contingency

10–25 % (Pre-PE): 25 %–60 %

25–35 % (Schematic): 20 %–40 %

35–70 % (Design Dev): 14 %–30 %

70–100 % (Final): 7 %–20 %

Editing tips:

Add or remove rows using the + Add band button and the ✕ icon at the end of each row. The table is free-form — it can have any number of rows as long as the first band starts at or below 1 and the last band extends to at least 100.

Labels are displayed in the envelope panel's tooltip, so use short, descriptive text (e.g. "5 % (SR-520 methodology)").

An agency with a tighter standard might tighten the late-stage band to 5 %–15 %; an agency with a fixed-price design-build preference might tighten earlier bands. The Restore defaults button reverts to the NHI values.

Warning: the envelope panel reads the project's Design Completion from Project Information. If that field is left blank, the envelope is not computed and the panel shows no state. Either set the completion percentage or leave the panel hidden by not clicking into the Triage view.

Table improvements (v1.10): the Design Completion table has been refreshed for clearer editing. The "% Design Completion" label column is now on the far left so users read row identity before numeric ranges. The completion-% columns (low/high) are back-shaded light blue and the contingency-% columns (low/high) back-shaded light orange to visually distinguish the two pairs. The delete "X" at the end of each row has been replaced with the standard trash icon used throughout the application and now prompts a confirmation dialog naming the band being removed. The "Restore defaults" button is now also enabled whenever the table has unsaved edits, not only when the server state differs from defaults - so it serves as a rescue path for accidental row deletions even when the project was at defaults moments ago.

11.6 How the Four Tables Interact

The four tables are edited independently but interact at runtime:

Cost scoring table → feeds the Cost distributions of individual risks AND the implied whole-project contingency used by the envelope panel. Raising Cost bands raises the envelope check's left side.

Schedule scoring table → feeds only the Schedule distributions of individual risks. It does not touch the envelope check.

Design Completion bands → feed only the envelope check's right side (the acceptable range). Raising the band highs relaxes the check; lowering them tightens it.

Probability Bands → feed P(occur) seeding when a workshop facilitator scores Prob in Triage and the row's P(occur) is still empty (see §11.4). They do not affect the envelope check or the cost / schedule distributions directly — they only translate the qualitative 1–5 Prob score into a numeric probability for the simulation.

Tip — if the envelope frequently flags Approaching or Exceeded on realistic projects, the likely causes (in priority order) are: (1) the Cost scoring table is too aggressive for your agency's project type; (2) the Design Completion contingency_high_pct values are too restrictive for the project's design stage. Adjust whichever matches your diagnosis and re-check against a sample of historical projects before settling on the new values.

Warning: restoring defaults on one table does NOT restore the others. Each table is reset independently. If you want to return the entire scoring system to defaults, click Restore defaults on all four tables.

Chapter 12: Distribution Health, Variance Share & What-If Sensitivity (added in v2.3 / v2.4)

Releases v2.3 and v2.4 added a coordinated set of analytical tools that help you decide whether a simulated CASTR result is defensible before you stake a budget on it. The improvements span three pages of the application:

• Risk Identification (Triage view) — the Contingency Envelope card was extended with SME-Reviewed columns, sortable headers, low-side warning flags, and a risk-count breakdown that shows reviewed vs unreviewed risks per phase-segment.
• Results — Summary tab — a new Distribution Health card gives you four lenses on the simulated output: P10–P90 spread, skewness, excess kurtosis, and three sub-sections covering variance concentration, risk-register coverage, and contingency by component.
• Results — Tornado tab — every Risk Register row on the YOE-cost tornado now carries a “% var” annotation showing its share of total variance, alongside the existing Spearman rank correlation. The Risk Impact view gained a Sort: EV / Sort: Variance Share toggle.

• Results — What-If Sensitivity tab (new) — an interactive panel that lets you deactivate one or more risks and watch the YOE distribution + contingency update live, without re-running Monte Carlo.
This chapter walks each feature in order. Definitions for the new terminology (variance share, SME-reviewed, percentage points, dominated / concentrated / broadly sourced bands, etc.) are in Appendix A: Glossary.

12.1 Triage Contingency Envelope — SME-Reviewed Coverage (v2.3.2)

On the Risk Identification page in Triage view, the Contingency Envelope card was extended with a three-way risk count column (Active / Reviewed / Inactive), two new low-side State flags, and sortable headers across every column. v2.5 split the SME-Reviewed columns out of this card: the Triage envelope now focuses on Seeded values (Seeded EV, Seeded Implied %) — the seed-stage view, where every priced risk counts toward the totals — while the defensible / sign-off-ready SME-Reviewed values live on a new SME-Reviewed Contingency Envelope card on the Risk Identification → Standard view (§5.1.6) and on the Risk Register (§5.2). The Triage card retains a single project-level SME-Reviewed Implied chip in its header so the workshop can see overall ratification progress without leaving the page.

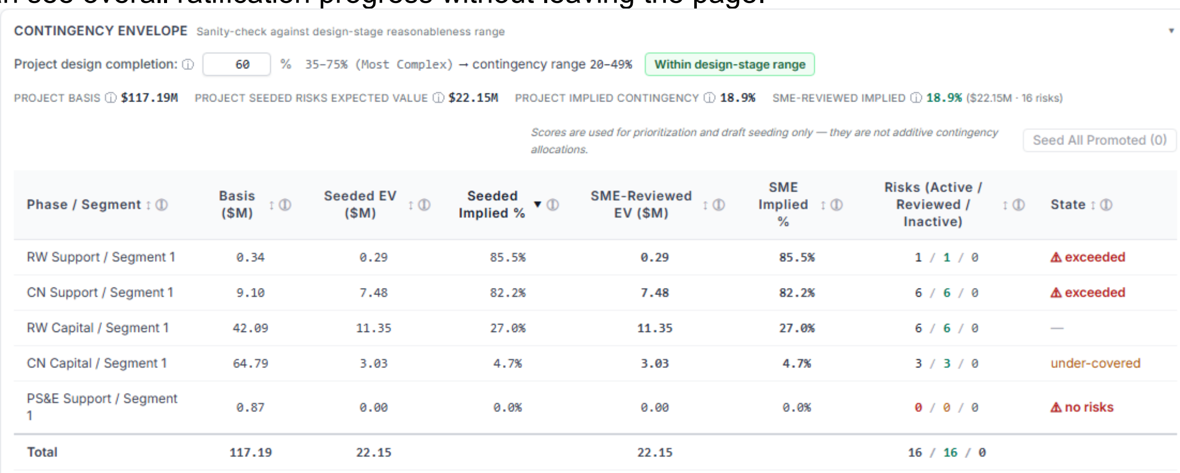
CONTINGENCY ENVELOPE Sanity-check against design-stage reasonableness range

Project design completion: 60 % 35-75% (Most Complex) → contingency range 20-49% Within design-stage range

PROJECT BASIS $117.19M PROJECT SEEDED RISKS EXPECTED VALUE $22.15M PROJECT IMPLIED CONTINGENCY 18.9% SME-REVIEWED IMPLIED 18.9% ($22.15M · 16 risks)

Scores are used for prioritization and draft seeding only — they are not additive contingency allocations. Seed All Promoted (0)

Phase / Segment	Basis ($M)	Seeded EV ($M)	Seeded Implied %	SME-Reviewed EV ($M)	SME Implied %	Risks (Active / Reviewed / Inactive)	State
RW Support / Segment 1	0.34	0.29	85.5%	0.29	85.5%	1 / 1 / 0	⚠ exceeded
CN Support / Segment 1	9.10	7.48	82.2%	7.48	82.2%	6 / 6 / 0	⚠ exceeded
RW Capital / Segment 1	42.09	11.35	27.0%	11.35	27.0%	6 / 6 / 0	—
CN Capital / Segment 1	64.79	3.03	4.7%	3.03	4.7%	3 / 3 / 0	under-covered
PS&E Support / Segment 1	0.87	0.00	0.0%	0.00	0.0%	0 / 0 / 0	⚠ no risks
Total	117.19	22.15		22.15		16 / 16 / 0	

Figure 12.1: Triage Contingency Envelope (v2.5 layout). Six sortable columns — Phase / Segment, Basis, Seeded EV, Seeded Implied %, Risks (Active / Reviewed / Inactive), and State. The SME-Reviewed columns moved to the Standard-view and Risk Register envelope cards in v2.5. Project Wide rows always pin to the bottom regardless of sort.

12.1.1 What “SME-Reviewed” means

The "SME-Reviewed" concept governs which risks count toward the SME-Reviewed Contingency Envelope cards on the Risk Identification → Standard view (§5.1.6) and on the Risk Register (§5.2), and toward the project-level SME-Reviewed Implied chip in the Triage envelope card's header. A risk's cost-side EV qualifies as SME-reviewed when ALL three conditions hold:
• The risk has a cost distribution seeded (cost_p1 is set).
• The cost-distribution “unconfirmed” flag has been cleared by opening the Distribution Editor and clicking Save (even with no edits — that click is the “an analyst has eyeballed and accepted these values” signal).
• The probability “unconfirmed” flag has been cleared by accepting the seeded P(occur) (clicking Enter in the field, or typing a different value).

Until both the cost-dist and probability flags are cleared, the risk is still on starter math; its EV contributes to Seeded EV but NOT to SME-Reviewed EV. The gap between the two columns is the work the workshop still has to do.

12.1.2 The new State flags: NO RISKS and UNDER-COVERED

Earlier versions of the State column flagged only high-side conditions (approaching / exceeded — the implied contingency is approaching or above the design-stage band). v2.3.1 added two low-side flags:

- ⚠ NO RISKS — zero risks of any kind are assigned to that phase-segment. This is the most common cause of an artificially skinny simulated distribution: with no inputs to draw from, the Monte Carlo simply has nothing to produce variance with.
- UNDER-COVERED — the implied contingency is below half the design-band's lower bound (e.g., band is 20–49 %, implied is < 10 %). This usually means the workshop has missed risks for that phase-segment OR set impact ranges too tight.

Use the Risks (Active / Reviewed / Inactive) count column alongside these flags to decide whether you need MORE risks or just need to widen the ranges on the ones you have. A 0 / 0 / 0 count with NO RISKS is unambiguous (add risks). A 5 / 5 / 0 count with UNDER-COVERED says “all risks reviewed but combined EV is small” — review the impact distributions.

12.1.3 Sortable headers and quick-edit dropdown

Every column on the Triage envelope card is sortable — click a header to sort, click again to flip direction. Sensible defaults apply: numeric and state columns default to descending (worst offenders first); the Phase / Segment text column defaults to ascending. Project Wide rows pin to the bottom of the table regardless of the active sort. The SME-Reviewed envelope cards on the Standard view (§5.1.6) and Risk Register (§5.2) are read-only summaries and are not click-sortable.

The # column on the main triage table itself was also made sortable (v2.3.2) for consistency with the other sortable triage columns (Name, Category, Phase / Segment, Priority, Quantify).

The Phase / Segment dropdown on the triage table now commits its value the moment you pick an option — you no longer have to click away or press Tab. This is intentional for selects since there is no “intermediate typing” state. Numeric and text fields elsewhere on the row still save on blur (so you can keep typing).

12.2 Distribution Health Card (v2.3.0, refined in v2.4)

On the Summary Statistics tab of the Results page, a new Distribution Health card sits above the existing YOE / CY StatCards. It answers the question “is this simulated distribution defensible?” without requiring statistical literacy from the reader.

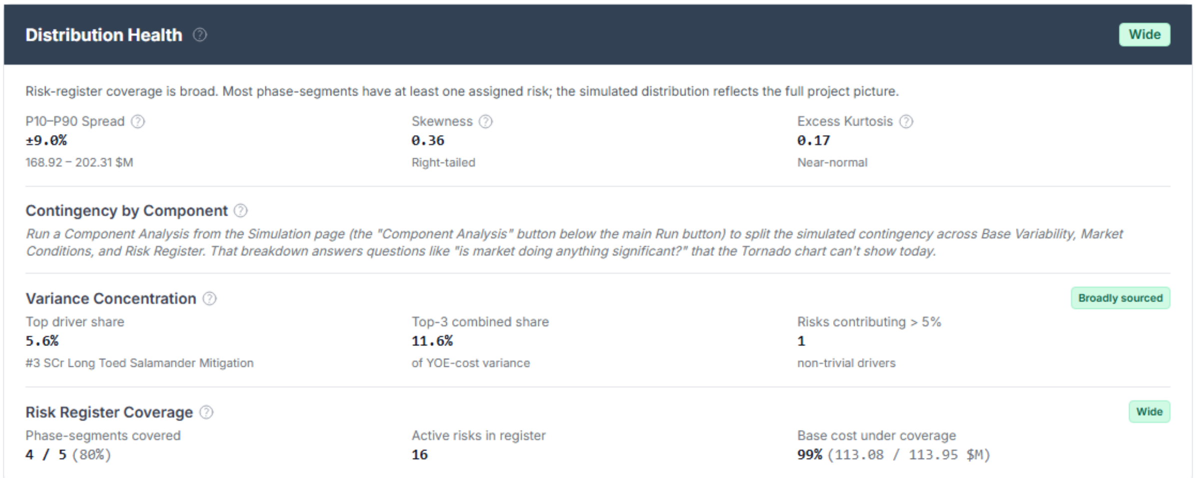

Figure 12.2: Distribution Health card showing the overall traffic-light, the three diagnostic cells (P10–P90 Spread, Skewness, Excess Kurtosis), the Variance Concentration block, and the Risk Register Coverage block.

12.2.1 The overall traffic-light

The Narrow / Moderate / Wide chip in the card header is driven by Risk Register Coverage alone — the actionable signal that drives the headline verdict. (The earlier v2.3.0 implementation used Coefficient of Variation to drive the chip; that benchmark band wasn't tied to the project's design completion stage and tended to mis-flag early-stage projects, so v2.4 removed it.) Hover the chip for a band-specific explanation:

• Narrow (red) — fewer than 50 % of phase-segments have any active risks. The simulated distribution is artificially tight; contingency is likely understated. Add risks to under-covered phase-segments first.

• Moderate (amber) — 50–80 % of phase-segments are covered. Result is usable; review the Coverage table at the bottom of the card to see which segments are still empty.

• Wide (green) — ≥ 80 % of phase-segments have an active risk in the register. The simulation reflects the full project picture and the result is defensible from a coverage standpoint.

12.2.2 The three diagnostic cells

Below the verdict, three cells provide the diagnostic detail:

• P10–P90 Spread — the practical 80 % confidence band, expressed as ± percent of the mean. This is the budget range cost engineers are usually asked to defend (“P10 to P90 is roughly $X to $Y, or a ± N% spread around the mean”).

• Skewness — measures asymmetry of the simulated YOE total cost distribution. Positive skew (right-tailed) is the natural shape of “cost can blow up but only go down so far.” Expect mild positive skew (~0.3 to ~1.5) on a typical project. See §12.5 below for interpretation rules.

• Excess Kurtosis — measures tail heaviness relative to a normal distribution. 0 = normal-shaped tails; > 0 = heavier tails (fat tails); < 0 = thinner tails. Independent of skewness — kurtosis is about tail mass, not asymmetry. See §12.5 below.

All three cells have inline help icons (ⓘ). The Skewness and Kurtosis icons give a brief definition; for the full “how to use it” interpretation, click the Guide button at the top of the page.

12.3 Variance Concentration block (v2.4.1)

Inside the Distribution Health card, the Variance Concentration block answers the workshop question “is one risk doing 80 % of the work?” with three numbers and a traffic-light chip:

• Top driver share — the share of YOE-total variance attributable to the single largest contributor. Shown alongside the risk's name (e.g., “#3 SCr Long Toed Salamander Mitigation”).
• Top-3 combined share — sum of the top three contributors' shares.
• Risks contributing > 5 % — count of risks with non-trivial variance contribution. A small count means the variance is concentrated in a few drivers; a large count means it's broadly sourced.

12.3.1 The traffic-light chip

• Dominated (red) — top driver > 50 %. One risk drives more than half of the YOE-cost variance. The Monte Carlo result is highly sensitive to that single risk's probability and impact distribution. Open the Tornado tab and sanity-check the dominant risk before defending the contingency.
• Concentrated (amber) — top driver > 30 % OR top-3 > 80 %. A small number of risks account for most of the variance. The result is usable, but the workshop should know which risks are doing the work — Tornado has the full ranking with each risk's exact % var.
• Broadly sourced (green) — top driver ≤ 30 % AND top-3 ≤ 80 %. No single risk dominates. Variance is distributed across many contributors via Central Limit Theorem aggregation, the most defensible posture.

12.3.2 The math

For each risk k, the variance share is computed exactly from the per-iteration impact matrix as Var(impacts_k) ÷ Var(total_yoe). The math is exact when risks are independent (Var of sum = sum of Vars). With copula-correlated risks the diagonal share is slightly inflated by cross-covariance terms, but the workshop verdict (“does any one risk dominate?”) is unchanged at the |ρ| levels typical for CASTR projects.

12.4 Contingency by Component strip (v2.4.3)

A three-cell strip showing Base Variability / Market Conditions / Risk Register dollars and % of total contingency. Mirrors the full Component Analysis table on the CASTR Contingency Analysis card; surfaced here so the user can answer “is market doing anything significant?” without scrolling.
This section is populated only after running a Component Analysis from the Simulation page — that sub-flow runs four sequential simulations to isolate each component's contribution; the regular single-run does not produce this breakdown. Until you run a Component Analysis, the section shows a guidance placeholder pointing the user to the right button.
The Market Conditions cell is colour-tinted by share so the answer is immediate visually:
• Gray (< 5 %) — market is negligible; the high “Base Cost” rows on the Spearman tornado are genuinely about base variability.
• Default (5–30 %) — market is contributing meaningfully but isn't a dominant driver.
• Amber (≥ 30 %) — market is a significant driver; the base-cost Spearman r values on the Tornado are partly market-inflated (because market multiplies base cost at the phase-segment level, so the rank correlation picks up the combined effect).
Why this strip is the canonical place to check market contribution: market conditions are sampled internally by the engine and applied as a multiplier to base cost — they don't appear as their own row on the Tornado. The strip is the only direct view of market's dollar contribution in the application today (a planned enhancement is to add market and base-cost variability to the per-iteration impact matrix, which would let them show up on the Tornado with proper variance shares).

12.5 Reading Skewness and Excess Kurtosis (v2.4)

Skewness and excess kurtosis are independent diagnostics of the simulated YOE distribution shape. They measure different things and answer different questions; both should be checked when defending a contingency.

12.5.1 Skewness — what to expect

Skewness measures asymmetry. For a typical CASTR project you should expect mild positive skew (~0.3 to ~1.5) — that's the natural shape of “cost can blow up but only go down so far.” Two readings worth investigating:
• Near-zero skew (|skew| < 0.3) on a project with discrete threats often means the threats are cancelling against opportunities, or the impact ranges are too symmetric. Review the risk register for unrealistic Triangle distributions.
• Very high skew (> 2.0) usually means one or two risks with very wide upper tails are stretching the distribution. Cross-check the Tornado tab to identify which, then sanity-check those risks' Mode/High values with the workshop.

12.5.2 Excess kurtosis — your fragility check

Excess kurtosis measures tail heaviness relative to a normal distribution. It is your “is the result fragile?” check.
• Excess kurtosis above ~1.0 (heavy tails) on a YOE total usually means one or a small number of risks are dominating the variance — the Monte Carlo result is sensitive to those few risks' parameters. Open the Tornado tab and the Variance Concentration block to confirm; if one risk is driving > 50 % of variance, sanity-check it before defending the contingency.
• Excess kurtosis near 0 (mesokurtic) is the comfortable case: variance is broadly sourced via Central Limit Theorem aggregation across many small contributors.
• Negative excess kurtosis (thin tails) is unusual on a real project register — it often signals missing low-probability / high-impact risks that the workshop hasn't surfaced yet, since real cost distributions almost always carry some tail weight.

12.5.3 Why CV (Coefficient of Variation) is no longer used

Earlier v2.3.0 used CV (std dev / mean) as the headline width metric with a hard-coded 10–30 % benchmark band. v2.4 removed CV because the benchmark wasn't tied to the project's design completion stage — the same CV value is interpreted very differently for an early-phase Class 5 estimate vs a 70 %-design Class 2 estimate. The fix is to tie CV bands to the existing design-completion contingency tables (a planned enhancement); until that's in place, CV is more confusing than helpful. The P10–P90 spread cell is a more directly defensible width metric (it's the practical 80 % confidence band cost engineers are usually asked to justify).

12.6 Variance Share on the Tornado (v2.4.1)

The Spearman tornado now carries a “% var” annotation on each Risk Register row. The bar length still shows Spearman rank correlation (which inputs matter most); the annotation shows the share of YOE-total variance attributable to that risk's per-iteration impact.

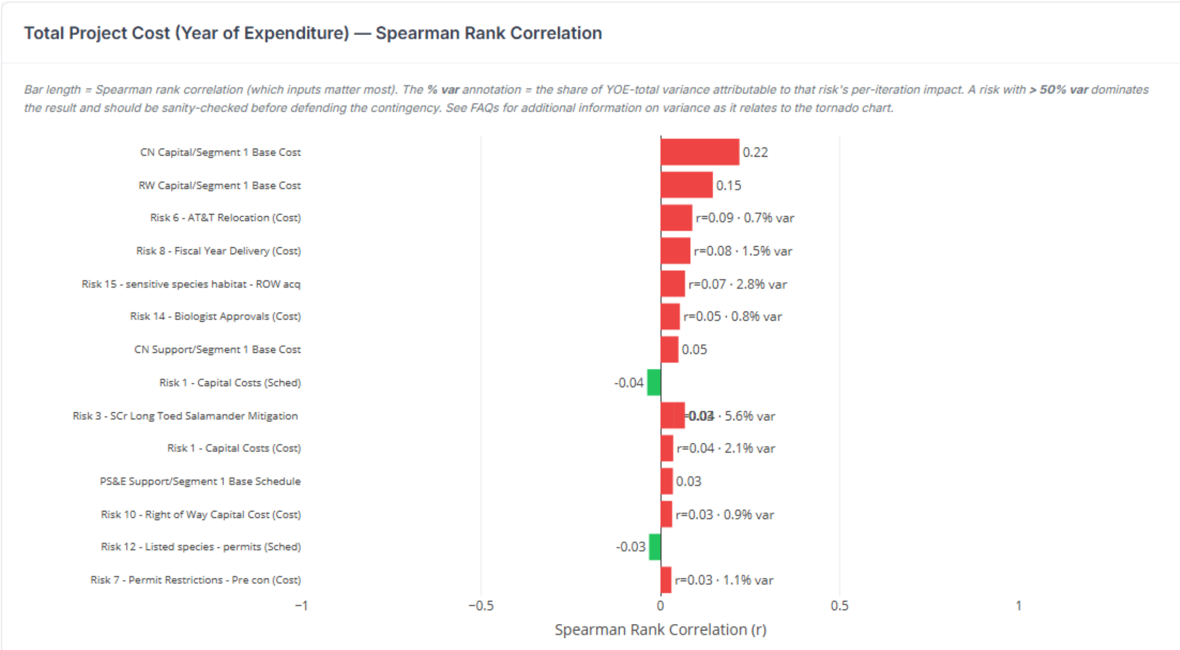

Figure 12.3: YOE Spearman tornado showing the new "% var" annotation on Risk Register rows. Base-cost rows and schedule-axis entries show only the Spearman r value (see why below).

12.6.1 Why some rows have no "% var" annotation

Three reasons, and they're not all the same:

• Base-cost variability rows and market condition multipliers contribute to total variance the same way risks do, but CASTR doesn't capture their per-iteration contributions in the impact matrix today (a planned enhancement). For now their share is implicit — if the displayed risk shares add to e.g. 65 %, the remaining 35 % is split between base, market, and risk-correlation cross terms.

• Schedule-axis entries can't be assigned a clean variance share at all because schedule risks aggregate via parallel max() in the engine and feed back into YOE cost through the inflation midpoint. Removing one schedule risk's contribution isn't a linear subtraction, so the leave-one-out math doesn't hold.

• Current Year (CY) tornado entries don't carry % var because the impact matrix is YOE-only by design.

12.6.2 Why CN/Seg X Base Cost rows often dominate the top of the chart

It's normal for base-cost variability rows to outrank Risk Register entries on the Spearman tornado — especially on Construction-heavy projects with mature estimates. The reason is structural: base-cost rows fire on EVERY iteration (probability = 1.0), while risks fire only at their P(occur). So a moderately-variable construction segment with a large dollar basis will routinely outrank a 30 %-probability risk whose actual EV is comparable.

This is normal and expected on the YOE Spearman tornado. The implication for contingency: if construction base-cost variability is driving most of the variance, the workshop should ask whether the Cost Var % on those rows is realistic for the project's design completion stage — narrower at 70 % design, wider at concept stage. The Risk Register is doing less work in that case; tightening the base-cost ranges (or, in the other direction, validating that they reflect real estimating uncertainty) is where the biggest contingency lever is.

12.6.3 Risk Impact tornado — Sort: EV / Sort: Variance Share toggle

On the Risk Impact view, every row name now also carries a “(X.X% var)” annotation, and the card header offers a Sort: EV / Sort: Variance Share toggle. EV (default) ranks by absolute expected value; Variance Share ranks by actual contribution to YOE-cost variance. The two orders frequently differ — a high-probability moderate-impact risk dominates EV but may not drive much variance, while a low-probability wide-range risk often does the opposite.

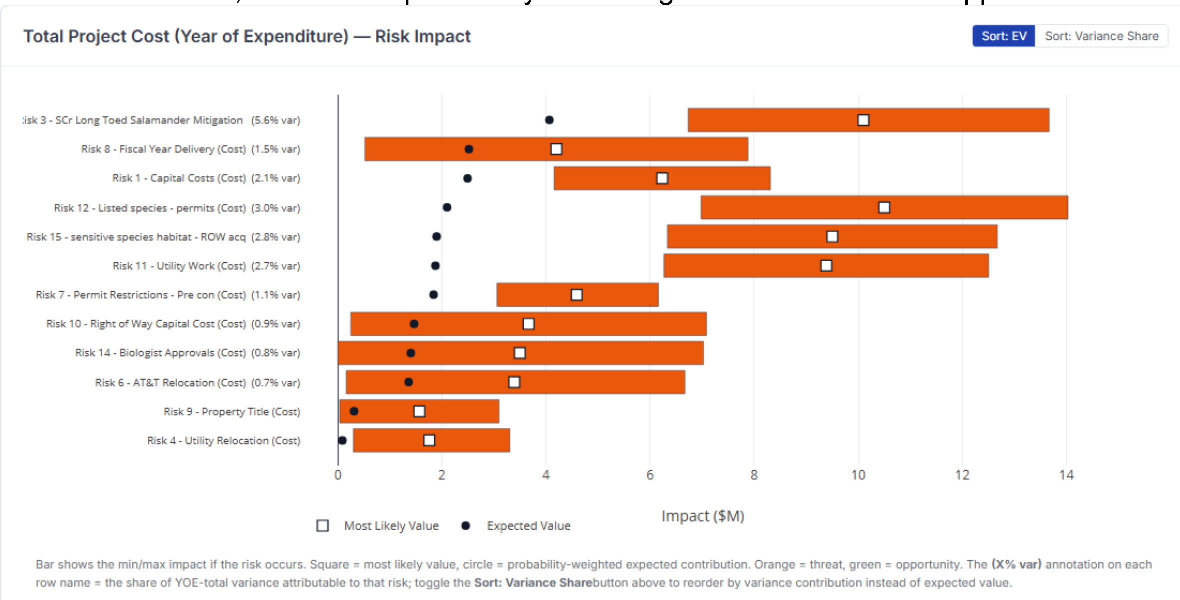

Figure 12.4: Risk Impact tornado with “(X.X% var)” annotations and the Sort: EV / Sort: Variance Share toggle in the card header.

12.7 What-If Sensitivity Tab (v2.4.0)

A new tab on the Results page that lets you deactivate one or more risks and watch the YOE distribution + CASTR contingency update in real time. Designed for workshop use — when a stakeholder asks “what if Risk #5 doesn't happen?” or “which two risks together are driving our overrun?”, you can answer immediately without re-running Monte Carlo.

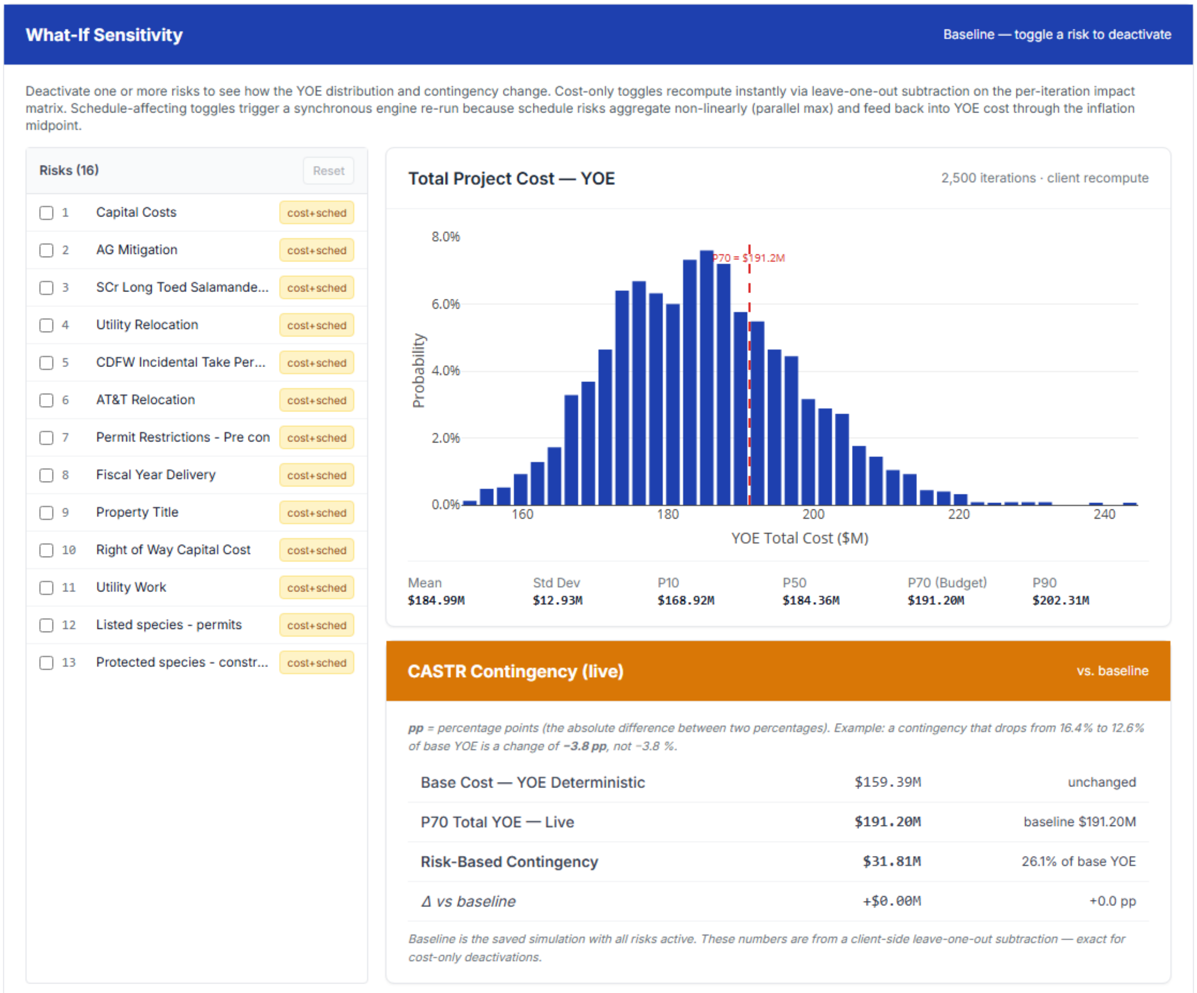

Figure 12.5: What-If Sensitivity tab in its baseline state, with all 16 risks active.

12.7.1 Layout

Three regions:

• Risk checklist (left) — every active risk in the register with its axis chip (cost, sched, or cost+sched). Schedule-affecting risks are tinted amber; cost-only risks are tinted green. The Reset button clears all deactivations in one click.

• Live YOE distribution chart (top right) — a histogram with a P70 marker that updates as risks are toggled. A six-cell stats strip below the chart shows Mean / Std Dev / P10 / P50 / P-budget / P90.

• CASTR Contingency (live) card (bottom right) — a slim version of the Summary tab's Contingency card, with a Δ vs baseline row in proper +/− formatting (the “pp” unit is explained in a small note above the table).

12.7.2 Hybrid recompute logic

Toggling a risk produces an immediate update via one of two paths, depending on the risk's axis:

• Cost-only deactivations recompute INSTANTLY in the browser via leave-one-out subtraction on the per-iteration impact matrix — the math is exact for cost variance. The badge under the chart title reads “client recompute”.

• Schedule-affecting deactivations (any risk with a schedule component) cannot be decomposed linearly because schedule risks aggregate via parallel max() and feed back into YOE cost through the inflation midpoint. The panel falls back to a synchronous engine re-run via the /whatif endpoint; the badge reads “engine re-run”. These take 1–2 seconds on small projects.

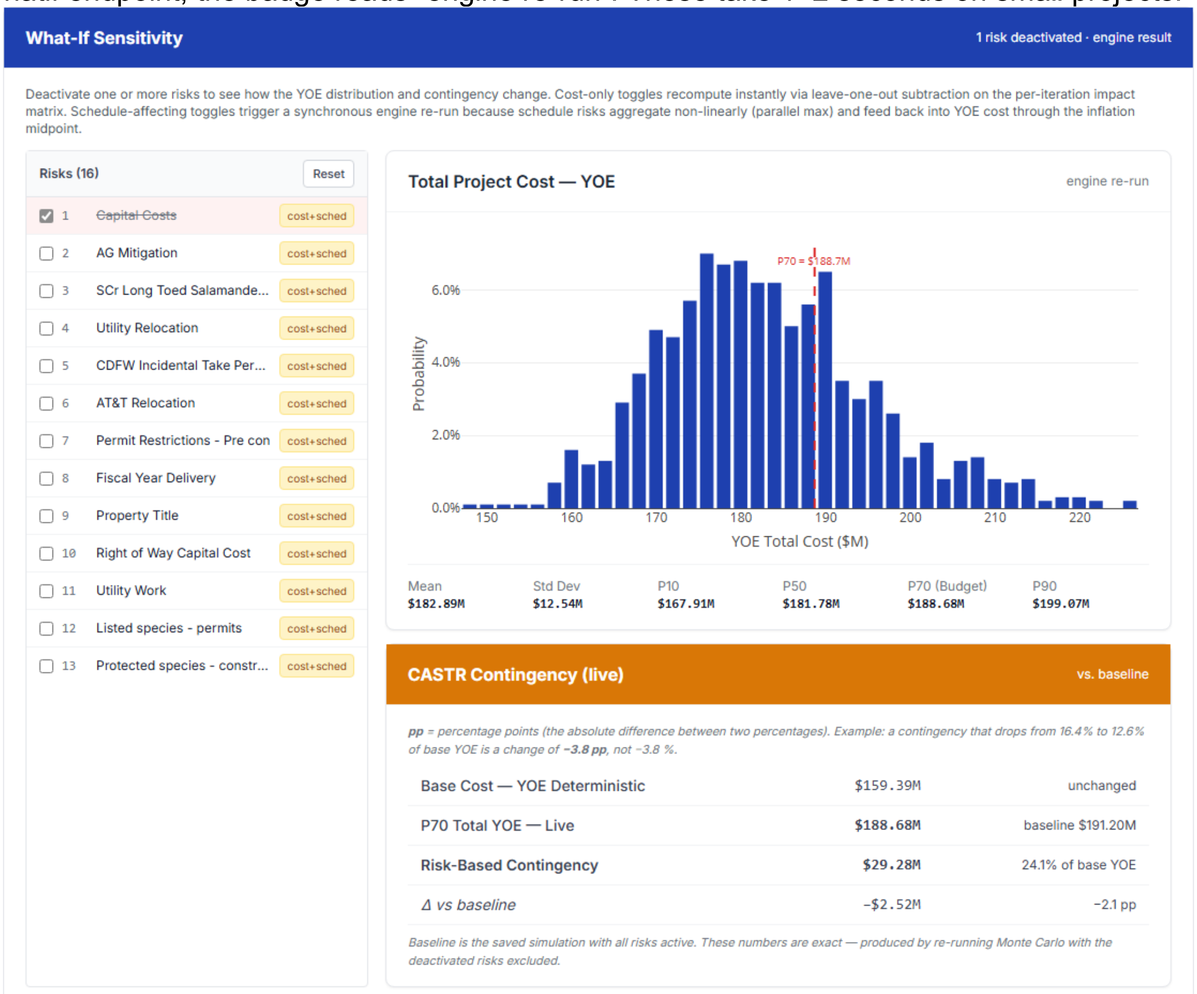

Figure 12.6: What-If Sensitivity tab after deactivating Risk #1 (Capital Costs, a cost+sched risk). The engine re-run badge under the chart title shows the path; the contingency strip's Δ vs baseline row reads −2.52M / −2.1 pp.

12.7.3 Workshop use vs. memo / board-pack use

Multi-risk deactivations involving correlated risks are slightly approximate when run via the client-side path (the Gaussian copula draws can't be perfectly undone by subtraction). For workshop discussion the approximation is fine — the verdict (“does contingency drop meaningfully?”) is robust. For a number that's going into a memo or a board pack, re-run the simulation from the Simulation page with the relevant risks set to include = False on the Risk Register, so you get a full saved run with exact numbers.

12.7.4 Reading small Δ vs baseline values

If the Δ vs baseline is materially smaller than the baseline's convergence wobble (Δ < ~0.5 % of mean YOE), it's within Monte Carlo noise. The Distribution Health and What-If panels both use the same iteration count as the saved run — if convergence was tight (use_convergence on, converged at modest iteration count), the noise floor is small. If the run was cancelled early or used a small iteration count, the floor is wider. The “iterations used” number under the chart title is the relevant context.

Appendix A: Glossary

The following table defines key terms used throughout CASTR and this manual.

Term	Definition
Base Cost	Neutral, deterministic current-year project cost excluding contingency and inflation, but including allowances for known but not fully defined scope items.
BetaPERT	A Beta distribution parameterized by minimum, most likely, and maximum values. Places more weight on the most likely value than a Triangular distribution, producing smoother tails.
BTP / AP / WTP	Better-Than-Planned, As-Planned, and Worse-Than-Planned market condition scenarios used to model bid environment.
Calibration Factor	A multiplier applied to the raw starter distribution values produced by the Seed action. Computed as clamp(this_phase_segment_Base_Cost / project_mean_phase_segment_Base_Cost, 0.75, 1.25). A factor of 1.00 leaves the raw values unchanged; 0.75 or 1.25 indicates the floor or ceiling of the clamp has been reached because the phase-segment is unusually small or large relative to the project average. The goal of calibration is to keep starter values sensible across phase-segments of very different sizes without letting the correction dominate the user's triage judgment.
CDF	Cumulative Distribution Function. The probability that a value is at or below a given threshold. Ranges from 0 to 1.
Checkpoint	A named, point-in-time snapshot of a project's Risk Identification and Risk Register tables. Created manually from the Checkpoints dialog or automatically before destructive operations (import, clear-all, restore, bulk delete). Used to support iterative workshop sessions by allowing the user to roll back to any prior state. See §8.4.
Clamp	A numeric operation that pins a value to a fixed range: clamp(x, low, high) returns low if x < low, high if x > high, and x otherwise. CASTR uses a clamp on the calibration factor at [0.75, 1.25] so the size correction applied to starter distributions can never multiply the raw values by less than 0.75 or more than 1.25, regardless of how unusual the phase-segment size is.
Combined Priority Score	An automatic 1–25 score computed in the Triage View as Probability × max(Cost, Schedule), using the risk's 1–5 triage scores. Higher values indicate risks that deserve the most attention during a quantification workshop.

Contingency	An amount added to the base estimate to account for project risk. In traditional estimating, contingency is typically historical-based and is derived from historical records of project performance, estimator experience, agency policy, or industry guidance for similar projects. In CASTR, contingency is developed through explicit quantitative risk modeling of base variability, market conditions, and identified risk events. In the CASTR application, the reported Risk-Based Contingency is the difference between the selected budget-percentile cost and the deterministic base cost.
Contingency Envelope	An industry-standard range of total project contingency (low % to high %) considered reasonable for a project at a given Design Completion percentage. CASTR compares the project's currently-implied whole-project contingency (derived from seeded cost distributions and their probabilities) against this envelope on the Triage view's envelope panel and flags Approaching or Exceeded states. The envelope bands default to the NHI 134205 values and are editable on Settings → Scoring Tables.
Convergence	The early stopping of a simulation when key statistics have stabilized within a specified tolerance, indicating that additional iterations would not meaningfully change the results.
Correlation	A statistical relationship between two uncertain variables in the risk model that describes how their sampled values move together during simulation. Correlation coefficients range from -1 to 1, where positive values indicate variables tend to move in the same direction, negative values indicate they tend to move in opposite directions, and zero indicates no linear relationship. In CASTR, these correlations are defined in the Correlation Editor and are applied during simulation using a Gaussian copula.
Cost Estimate	The sponsor's full cost number for a project or phase-segment, including the Traditional Contingency line. On the Base Variability page this is the number entered in the "Cost Estimate (including Contingency)" column. Distinct from Base Cost, which is Cost Estimate minus the Traditional Contingency amount.
COV	Coefficient of Variation. The standard deviation divided by the mean, expressed as a ratio. Provides a normalized measure of dispersion.
Current Year (CY)	Costs expressed in today's dollars without any inflation adjustment. Also referred to as constant dollars or base-year dollars.
Design Completion %	An integer 0–100 set on the Project Information page indicating how far the project design has progressed. Drives the selection of the Contingency Envelope band used on the Triage view's envelope panel. Typical conventions: 1–10 % Concept, 10–25 % Pre-PE, 25–35 % Schematic, 35–70 % Design Development,

	70–100 % Final. The exact band labels and contingency ranges are configurable on Settings → Scoring Tables.
Expected Cost	The probability-weighted mean cost impact of a risk: Probability of Occurrence × mean of the Cost distribution. Used in portfolio totals, contingency envelope computation, and tornado-chart ranking. Threats have a positive Expected Cost; Opportunities have a negative Expected Cost.
Gaussian Copula	A mathematical method for modeling correlations between multiple random variables simultaneously while preserving their individual marginal distributions.
Heat Map	A 5 × 5 grid that counts risks by their Probability score (row) and Cost or Schedule score (column), colour-graded from light green (low severity) to deep red (upper-right high-severity corner). The Triage view shows two heat maps side by side — one for Cost, one for Schedule. Clicking cells filters the Triage table; multiple cells can be selected at once.
Mitigation	An intervention that reduces a risk's probability of occurrence and/or its cost or schedule impact. Entered per-risk through the Mitigation Details panel on Risk Identification. CASTR tracks the pre-mitigation (original) and post-mitigation (residual) distributions separately; the Run Mode on the Simulation page selects which is used. See Chapter 10.
Monte Carlo Simulation	A computational technique that uses repeated random sampling to model the probability distributions of possible outcomes for a system with inherent uncertainty.
PDF	Probability Density Function. Describes the relative likelihood of each possible value. The area under the curve between two values gives the probability of the outcome falling in that range.
Percentile (P-value)	The value below which a given percentage of observations fall. For example, P70 means 70% of simulation outcomes are at or below this value.
Phase-Segment	A specific combination of a project phase (e.g., Construction) and a geographic segment (e.g., Main Corridor). Phase-segments are the fundamental unit of analysis in CASTR.
Predecessor	A phase-segment that must complete (or reach a defined milestone) before another phase-segment can begin.
Quantify Flag	A per-risk toggle in the Triage view indicating whether the risk is intended to graduate to a full cost/schedule distribution (Quantify = yes) or remain qualitative / be handled on the Minor Risks list (Quantify = no). The Seed action only operates on risks flagged Quantify = yes.
Raw vs Calibrated Values	Two forms of a starter distribution produced by the Seed action. Raw values are the output of the scoring-table lookup applied directly to the phase-segment Base Cost (for cost) or directly converted from weeks to months (for schedule). Calibrated

	values are the raw values multiplied by the Calibration Factor. The calibrated form is what gets persisted to the risk; the raw form is shown in the Seed confirmation dialog's notes so the user can see how the persisted number was derived.
Residual Distribution	The post-mitigation cost or schedule distribution for a risk marked for Mitigation, representing the risk's expected behaviour if the mitigation strategy is implemented. Edited through the Risk Details panel alongside the pre-mitigation distributions.
Risk-Based Contingency	The portion of project funding above the deterministic base cost that is needed to achieve a selected confidence level, such as P70, in the simulation results. In the CASTR application, Risk-Based Contingency is reported as the difference between the selected budget-percentile cost and the deterministic base cost. This term is the software's reported contingency result and should be distinguished from the broader estimating concept of contingency.
Scoring Tables	Three editable global tables driving the Seed action and the Contingency Envelope check: the Cost Scoring Table (cost_score 1–5 → low/mode/high fractions of Base Cost), the Schedule Scoring Table (schedule_score 1–5 → low/mode/high weeks), and the Design Completion Bands (envelope ranges by design percentage). Edited on Settings → Scoring Tables. Changes apply only to subsequent Seed operations; already-seeded distributions are not updated retroactively. See Chapter 11.
Seed Distribution	A starter Triangular cost and/or schedule distribution generated automatically from a risk's 1–5 triage scores. Cost bands come from the Cost Scoring Table applied to the phase-segment Base Cost; schedule bands come from the Schedule Scoring Table converted from weeks to months. Both are then scaled by the Calibration Factor. Invoked per-risk via the Seed button on the Triage view or in bulk via Seed All. Seeded values are a starting point — the Distribution Editor can refine them.
Spearman Rank Correlation	A non-parametric sensitivity statistic used in tornado charts to rank how strongly an input variable is associated with total project cost or total schedule outcome based on rank order. It is an output diagnostic used for prioritization and sensitivity analysis, not a model input used to link variables during simulation. A high |Spearman rank correlation| indicates the variable strongly influences the simulated outcome's rank order. It is not the same thing as the input correlations defined in the Correlation Editor — those are a model input; the Spearman value is a model output.
Starter Library	An app-global, editable library of predefined risk categories and risk events shared across every project on this CASTR install. Accessed via Settings → Starter Library. Edits do not affect previously imported risks. See Chapter 9.

Tail Driver	A risk that is disproportionately present (or absent) in the extreme outcomes of a simulation, contributing significantly to best-case or worst-case scenarios.
Triage View	An alternative Risk Identification layout optimised for rapid 1–5 qualitative scoring of Probability, Cost, and Schedule. Paired with two heat maps, a Quantify flag, a Combined Priority column, and a Seed button that produces starter distributions from the triage scores. Intended for workshop-style sessions before committing to full quantification. See §5.1.2.
Triangular Distribution	A probability distribution defined by three parameters: minimum, most likely, and maximum values. The most commonly used distribution in risk analysis.
Uniform Distribution	A probability distribution where all values between the minimum and maximum are equally likely. Used when there is no basis for preferring one value over another.
Year of Expenditure (YOE)	Costs adjusted for inflation to reflect the actual dollar amounts at the time expenditures will occur. Accounts for the time value of money.
Distribution Health card	A diagnostic card on the Results page Summary Statistics tab that gives four lenses on the simulated YOE total cost: P10–P90 spread, skewness, excess kurtosis, and three sub-sections covering Variance Concentration, Risk Register Coverage, and Contingency by Component. The header traffic-light (Narrow / Moderate / Wide) reflects Risk Register Coverage.
Skewness	A measure of asymmetry of the simulated distribution. Positive skew (right-tailed) is the natural shape of cost distributions — there are more ways for projects to overrun than to come in under. Computed as the sample third-moment ratio $m_3 / m_2^{1.5}$. For a typical CASTR project, expect mild positive skew (~0.3 to ~1.5).
Excess kurtosis	A measure of tail heaviness of the simulated distribution relative to a normal distribution. Normal = 0; positive = heavier tails (fat tails); negative = thinner tails. Independent of skewness — kurtosis describes tail mass, not asymmetry. Computed as $m_4 / m_2^2 - 3$. Values above ~1.0 typically signal that one or a small number of risks are dominating the variance.
Variance share	The fraction of YOE-total variance attributable to a single risk's per-iteration impact, computed exactly from the per-iteration impact matrix as Var(impact_k) ÷ Var(total_yoe). Surfaced as a “% var” annotation on each Risk Register row of the Tornado chart and aggregated into the Variance Concentration block on the Distribution Health card. Exact for independent risks; mildly inflated by cross-covariance for correlated risks but the workshop verdict is unchanged.
Variance concentration	A summary diagnostic on the Distribution Health card showing how much of the YOE-cost variance is driven by a single dominant risk. Three numbers: top driver share, top-3 combined share, count of risks contributing > 5 %. The traffic-light chip reads Dominated (top

	driver > 50 %), Concentrated (top driver > 30 % or top-3 > 80 %), or Broadly sourced (otherwise).
SME-reviewed (risk)	A risk whose cost distribution AND probability have both been explicitly accepted by the workshop. A risk qualifies as SME-reviewed when (a) it has a seeded cost distribution, (b) the cost-dist seeded_unconfirmed flag has been cleared (Distribution Editor opened and saved), and (c) the probability seeded_unconfirmed flag has been cleared (P(occur) accepted via Enter or typed). Seeded EV in the Triage envelope card includes all seeded risks; SME-Reviewed EV includes only the subset where both flags have been cleared. The gap is the work the workshop still has to do.
NO RISKS / UNDER-COVERED (Triage State flags)	Two low-side warnings on the Risk Identification Triage envelope card. NO RISKS = zero risks of any kind are assigned to that phase-segment. UNDER-COVERED = the implied contingency for that phase-segment is below half the design-band's lower bound. Both indicate that the simulated distribution will be artificially tight unless the workshop adds risks (NO RISKS case) or widens existing impact ranges (UNDER-COVERED case).
What-If Sensitivity	A tab on the Results page that lets the user deactivate one or more risks and watch the YOE distribution + CASTR contingency update live. Cost-only deactivations recompute instantly via leave-one-out subtraction on the per-iteration impact matrix; schedule-affecting deactivations trigger a synchronous engine re-run via /whatif because schedule risks aggregate non-linearly (parallel max) and feed back into YOE cost through the inflation midpoint.
Per-iteration impact matrix	A two-dimensional array of shape (n_iter × n_risks) captured by the simulator at run time. Each cell contains a single risk's YOE-signed cost contribution to a single iteration (already gated by the Bernoulli draw and scaled by the YOE inflation factor). The matrix is persisted with the saved run and powers the What-If Sensitivity panel's instant client-side recompute, plus the variance share annotations on the Tornado and Distribution Health cards.
Percentage points (pp)	The absolute difference between two percentages. A contingency that drops from 16.4 % to 12.6 % of base YOE is a change of −3.8 pp — NOT −3.8 % (which would mean a relative drop of 3.8 % of 16.4 %, much smaller). Used anywhere two percentages are compared in CASTR (e.g., the Δ vs baseline row on the What-If Contingency strip) to avoid the relative-vs-absolute ambiguity. Standard convention in cost-engineering and finance reporting.
P10–P90 spread	The width between the P10 and P90 of the simulated YOE total cost distribution, expressed as ± percent of the mean. The practical 80 % confidence band cost engineers are usually asked to defend. A narrow spread (e.g., ± 5 %) means the simulation considers very few outcomes far from the mean; a wide spread (± 25 %+) reflects substantial uncertainty appropriate to early-stage estimates.
Coverage Warning Level	The State column on the Triage Contingency Envelope card combines high-side and low-side warnings. High-side: approaching (within 85–105 % of the design-band upper bound) and exceeded (above 105 %). Low-side (added in v2.3.1): under-covered (implied

	% below half the band lower bound) and no_risks (zero risks assigned to that phase-segment). High-side warnings take precedence when both apply.

Appendix B: Distribution Types

CASTR supports the following probability distribution types for modeling cost and schedule uncertainty. Select the appropriate distribution based on the nature of the risk and the quality of available information.

Triangular

The Triangular distribution is defined by three parameters: Minimum (P1), Most Likely (P2), and Maximum (P3). It is the most commonly used distribution in risk analysis because it is intuitive and easy to elicit from subject matter experts. The distribution forms a triangle shape with the peak at the most likely value.

Use the Triangular distribution when you can confidently estimate a minimum, most likely, and maximum value for the risk impact. It is appropriate for most risk events in highway construction projects.

BetaPERT

The BetaPERT distribution is parameterized by Minimum (P1), Most Likely (P2), and Maximum (P3), similar to the Triangular distribution. However, BetaPERT places more weight on the most likely value, producing a smoother, more bell-shaped curve with less emphasis on the extreme minimum and maximum values.

Use BetaPERT when you believe the most likely value is a strong central estimate and the extremes are less probable than a Triangular distribution would suggest. BetaPERT is recommended when expert judgment is well-calibrated.

Uniform

The Uniform distribution is defined by two parameters: Minimum (P1) and Maximum (P2). Every value within the range is equally likely. The distribution forms a flat rectangle between the minimum and maximum.

Use the Uniform distribution when there is no basis for preferring one value over another within a known range. This is common for early-stage estimates where uncertainty is high and no most likely value can be identified.

Normal

The Normal (Gaussian) distribution is defined by two parameters: Mean (P1) and Standard Deviation (P2). It produces the familiar symmetric bell curve centered on the mean value.

Use the Normal distribution when the risk impact is well-characterized by historical data or statistical analysis, and when positive and negative deviations from the mean are equally likely. Note that the Normal distribution has infinite tails, meaning extremely large or small values are theoretically possible, though with very low probability.

Constant

The Constant “distribution” (added in v1.9) is a degenerate case with a single parameter, Value (P1). Every simulation iteration returns the same Value with no sampling variation, which makes it the right choice when a risk impact is known exactly (a fixed fee, a scheduled contractual payment, a known acceleration credit). Its preview chart shows a vertical spike at the chosen value. Selecting Constant in the Distribution Editor collapses the parameter form to a single Value field.

Use Constant when you have high confidence in a single number and do not want variance from that input to propagate into the output distribution. For unknown costs or durations with any uncertainty, prefer Triangular, BetaPERT, Uniform, or Normal instead.

Parameter Validation

Starting with v1.9, every distribution is validated before the save is allowed. The rules are identical across the Distribution Editor, Risk Details residual editor, and the Market Conditions BTP / As-Planned / WTP triples. A failing validation blocks the save, keeps the editor open, and displays a red banner naming the exact rule that was violated.

Triangular requires Minimum ≤ Most Likely ≤ Maximum and Minimum < Maximum (ramp-shaped triangles with the mode at the boundary are allowed). BetaPERT is stricter: Minimum < Most Likely < Maximum (a boundary mode is rejected because the reparameterization degrades). Uniform requires Minimum < Maximum. Normal requires Standard Deviation > 0; any finite Mean is accepted. Constant requires a single finite Value — zero, negative, and very large values are all legitimate user choices.

Rows that are still being filled in — a distribution whose three parameters are not all entered yet, or a market-condition triple where only some of Low / Most Likely / High are set — are treated as “in progress” and are not validated until all required values are present. This lets you enter values in any order during workshops without being blocked mid-edit, and incomplete rows are safely skipped by the simulation engine rather than aborting the run.

Appendix C: Keyboard Shortcuts

The following keyboard shortcuts are available within CASTR to improve efficiency.

Shortcut	Description
Ctrl+V	Paste a column of values from the clipboard into a table. Click the first cell in the target column, then press Ctrl+V to paste. This supports direct paste from Excel spreadsheets.
Enter Key	In the Risk Identification Triage View quickly enter a new risk line by placing cursor into the risk event name field and click the enter key instead of clicking the Add Risk button. This is useful when brainstorming with a large group of experts to open up additional risk rows to capture audience feedback.