CASTR User Manual

Version 2.15.0 · Updated 2026-05-12

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:Files or C:).
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.

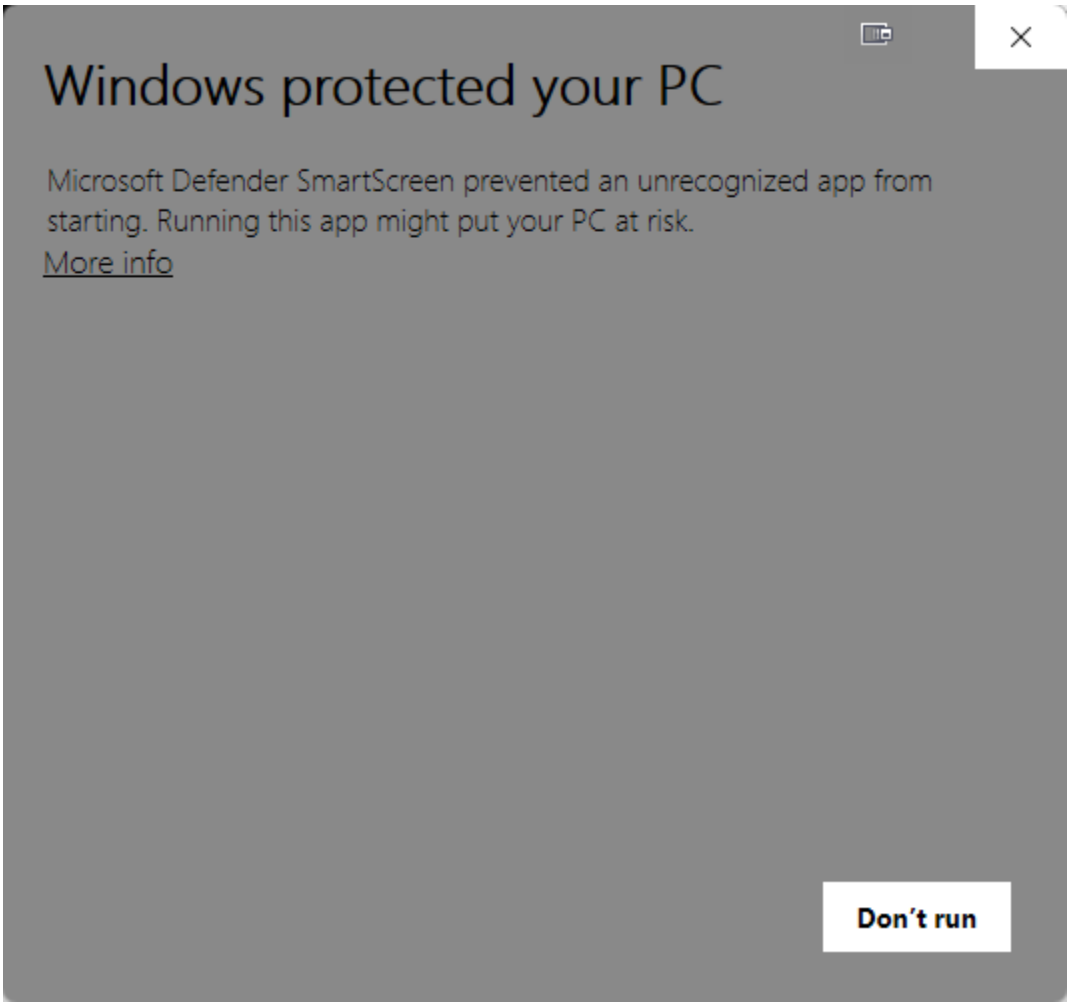

Windows SmartScreen may also display a warning because the application is not yet code-signed. If you see this message, click “More info” and then “Run anyway” to proceed.

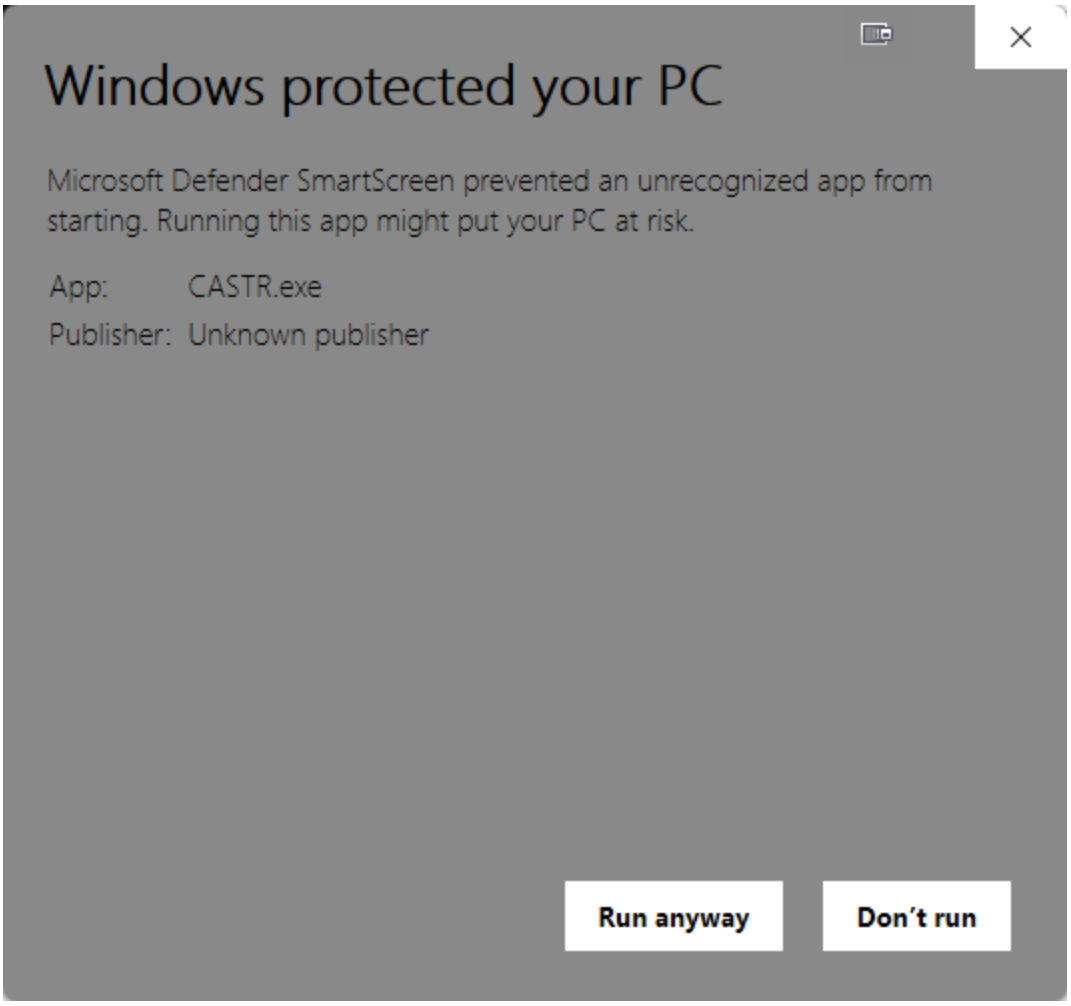

Updating CASTR

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).

Splash screen. On launch, CASTR shows a brief CASTR splash screen layered over the project list. Click **OK** in the lower-right corner (or press **Enter**, **Space**, or **Esc**) to dismiss it; the splash fades out smoothly to reveal the project list underneath. The splash is shown once per app launch — navigating between pages does not bring it back.

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.

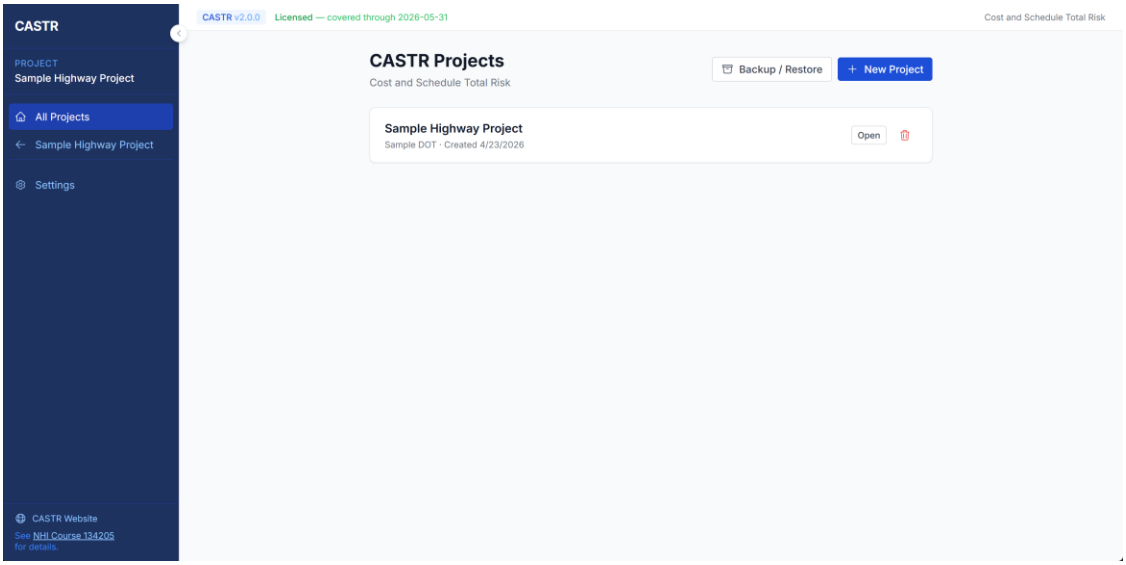

Figure 1.1 - Project List home page

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.

Notifications. Status messages and errors from background operations (saving a row, importing risks, exporting Word, etc.) appear as small toast notifications in the bottom-right corner of the screen. Successful operations show in green and auto-dismiss after about five seconds; errors show in red and stay a little longer so you have time to read the message. Toasts are non-blocking — you can keep working while one is visible, and you can hover over a toast to pause its dismissal timer.

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: 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.

Bundled demo project. New installations seed a **Rural Highway Widening (4 to 6 Lanes) - Demo** project (DEMO-RHW-26-001) on first run — a 28-mile rural arterial widening that exercises the full feature set, including PoS@P70, tornado and tail drivers, MI/ME risk dependencies, inflation sensitivity, cost↔schedule correlation, and per-axis response ROI. Existing projects in your database are not affected.

1.7 Troubleshooting

CASTR is designed to recover gracefully from unexpected errors so you never lose work. Two recovery surfaces are worth knowing about:

Recovery card. If a page encounters an unexpected error while rendering (for example, a temporary inconsistency between a saved simulation run and the current input data), CASTR replaces the page with a recovery card titled “Something went wrong.” Your data is safe — nothing has been saved or deleted. The card offers three actions: *Try again* re-

renders the page (useful when the error was transient, such as a query that's now finished loading); *Back to projects* returns you to the project list so you can re-enter the project from a clean state; *Reload page* reloads the entire CASTR app (last resort). The error message itself is shown in the card and is also written to the browser console for support. If you see this card and the issue persists, take a screenshot and email it to support along with the steps that led up to it.

Lost connection. If the CASTR background server stops responding (rare — typically only if the desktop window has been closed and reopened, or if a Windows update has interrupted it), the application shows a connection-down banner and retries automatically. Closing and reopening CASTR from the Start menu always restores normal operation.

1.8 Terminology — 'response' vs. 'mitigation'

CASTR uses a per-axis risk-response model. A single risk can carry different PMI response strategies on its cost and schedule axes — for example, Mitigate on the cost-threat side while Enhance is selected on the schedule-opportunity side. To accommodate both threat-side strategies (Mitigate, Avoid, Transfer, Escalate, Accept) and opportunity-side strategies (Exploit, Enhance, Share, Escalate, Accept) under one umbrella, the user-facing terminology is 'response' rather than 'mitigation'.

The user-facing labels read Response Analysis, Response What-If, Response ROI Summary, RES (the master toggle on the Risk Register), Response Cost, Response Owner, Response Due Date, and Pre-Response / Post-Response Run Modes. The PMI strategy names themselves are unchanged — Mitigate remains a threat-response option in the Risk Response dropdown alongside Avoid, Transfer, Accept, and Escalate, and Enhance remains the opportunity-side counterpart alongside Exploit, Share, Accept, and Escalate. Saved project files, exports, and API field names retain their original identifiers (mitigation_cost, mitigation_owner, mitigation_due_date, mit_active, etc.) so existing data loads without migration.

1.9 Recommended Workflow

CASTR is organized as a left-to-right, top-to-bottom checklist: every page in the sidebar feeds the next, and a page typically cannot produce useful output until the pages above it have been completed. Working through the sidebar in order is the fastest way to a defensible result, and it is also the order the **Progress Checklist** tracks. The checklist appears in the sidebar whenever a project is open and turns green item by item as you complete each step; clicking any incomplete item navigates directly to the relevant page.

The full project workflow is twelve pages, grouped into five phases. The first time through a real project, plan to spend most of your effort on the four risk pages (steps 6–9), since the quality of the results depends almost entirely on the quality of those inputs.

Setup — define what you're modeling.

1. **Project Info.** Open or create the project, fill in the basic metadata (name, owner, location, descriptions) and set the *Estimate Last Update Date*. That date anchors every inflation calculation downstream, so it must be set before you do anything else.
2. **Segments & Phases.** Build the work-breakdown structure. A project has one or more segments (geographic or contractual chunks); each segment has phases (PE, ROW, UTL, CN, CE, etc.). Enter the deterministic base cost and base duration for every phase — these are the anchors the Monte Carlo perturbs.
3. **Inflation Tables.** Configure the annual inflation rates that escalate base-year dollars to year-of-expenditure dollars. You can author multiple sensitivity scenarios on this page; the simulation lets you choose which scenario to run.

Inputs — describe the unknowns.

4. **Base Variability.** For each phase, enter the base-uncertainty distribution (typically a triangular low/most-likely/high) representing pure estimating uncertainty *before* any named risks fire. This is the cost spread the project would have if no specific risks materialized.
5. **Market Conditions.** Enter the bid-market premium distribution — how much the contracting environment is expected to inflate or deflate winning bids relative to the engineer's estimate.
6. **Review Staging.** Lay out the schedule: when each phase starts and how its spend is profiled across the years of expenditure. Combined with the Inflation Tables, this is what converts base-year dollars into YOE dollars.

Risks — identify, quantify, and link.

7. **Risk Identification.** Brainstorm the risk universe. Each entry gets a description, threat-or-opportunity classification, and qualitative L×I scores from the Scoring Tables. Risks here are inactive — they do not enter the simulation yet. The Starter Library (sidebar) seeds this page from a curated catalog you can edit. Promote the risks you want to model into the Risk Register using *Activate*.
8. **Risk Register.** The active risk list. For every risk add the quantitative fields the simulation needs: probability of occurrence, cost-impact distribution, schedule-impact distribution, predecessors (which phases the risk affects), MI/ME dependencies (mutually inclusive / mutually exclusive groups), and the response strategy with pre- and post-response distributions, response cost, owner, and due date.
9. **Correlation.** Set pairwise correlations between risks (and between risks and base-variability factors). These flow into the Gaussian copula that links variables during sampling. Defaults are zero; only edit pairs with a known relationship.
10. **Minor Risks.** Bulk-enter small risks too numerous to model individually. They are aggregated into a single contingency draw rather than getting their own rows in the register.

Analysis — run the model and read the answer.

11. **Run Simulation.** Choose the run mode (Pre-Response, Post-Response, or both), the iteration count (10,000 is the usual sweet spot), the inflation scenario, and an optional random seed for reproducibility. Click *Run* and watch the progress stream.
12. **Results.** Read off P50/P70/P80/P90 contingency values, the cost and schedule CDFs, the **Cost of Delay** card, the **Tornado** chart for variance share, **Tail Drivers** for what shows up in the worst 10% of outcomes, **Distribution Health** for input diagnostics, **Response What-If / ROI** for response-toggle analysis, **PoS@P70** for probability-of-success at a chosen budget percentile, and the **YOE Cashflow** chart for year-by-year spend. From this page you can export the full report to Excel, Word, or PDF.

Iteration is normal. A first pass through these twelve pages produces a baseline. Most projects then iterate — refining risk distributions, adjusting correlations, testing alternative responses, comparing inflation scenarios — and re-running the simulation as inputs are tightened. Each project supports duplication from the Project Info page (the *Duplicate…* button in the header opens a dialog for the new project's name) so you can run side-by-side scenarios without disturbing the baseline.

Where to get help while you work. Every page has a **Guide** button (a top-to-bottom walkthrough of that page) and an **FAQ** button (page-specific Q&A) on the right edge of the screen. Field-level question-mark icons open short tooltips with deep links into this manual.

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 "Duplicate..." button in the page header (top right). A dialog opens explaining what gets copied — segments, phases, risks, correlations, inflation tables, base variability, etc. — and prompting you for the new project's name. Click "Duplicate" to confirm; the new project opens automatically. Simulation history is intentionally not copied so the new project starts with a clean analysis.

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.

2.6 Reporting Errors and Generating Diagnostic Reports

If CASTR shows an error or behaves unexpectedly, send a diagnostic report along with your support email so the developer can reproduce and fix the issue without needing access to your machine.

There are two paths, depending on what you are looking at when the issue happens.

Path A — Copy details from a current simulation error

When a simulation fails on the Run Simulation page, the red Simulation Error pane has a **Copy details** button at the top right. Click it. Your clipboard now contains a single text block with:

- The CASTR application version
- The current URL and your browser
- Run Mode, Iteration count, the three component toggles (Base Variability, Market Conditions, Risk Register)
- Your inflation scenario list and Budget Percentile
- The full error message including the Python traceback

Paste this directly into your support email. This is the fastest path when you want to report the specific simulation that just failed.

Path B — Download a broader diagnostic report

When the issue is not a current simulation error (e.g. the YOE chart looks wrong, the Tornado tab is empty, the page failed to render), open **Settings** in the left sidebar and click **Download Diagnostic Report**. Your browser saves a JSON file named like `castr-diagnostic-YYYYMMDD.json`. Attach the file to your support email.

The JSON file contains:

- **Version info** — CASTR application version, release date, Python version, operating system, report generation timestamp.
- **Recent failed simulation runs** — the most recent ten errored runs across the entire database, each with its full Python traceback. (CASTR persists every failed run's traceback so you can attach old failures, not only the one currently on screen.)
- **Per-project shape summary** — for every project: number of phase-segments, number of risks (total / active / response-active), number of correlation pairs, and total simulation runs (with the count of errored runs). Project names are truncated to the first 60 characters.
- **Browser info** — user agent, viewport size, screen size, language, timezone, current URL.

What the diagnostic report does NOT contain

The report deliberately excludes:

- License key and license email
- Raw cost figures and project basis
- Full risk event names beyond the first 60 characters
- Simulation result data (cost distributions, percentiles, tornado entries)
- Risk distribution parameters

You can confidently attach the file to a support email without exposing project-sensitive data. The bundled metadata is enough for the developer to reproduce the issue against a synthetic project of the same shape.

Workflow summary

1. **Email a clear description** of what you did and what you saw versus what you expected.
2. **Attach the diagnostic report** (Path B) or **paste the copied details** (Path A).
3. **Optionally include a project backup ZIP** if the issue depends on specific data — Project list → Backup / Restore → Export. This is sensitive data; only include it if the developer asks.

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.

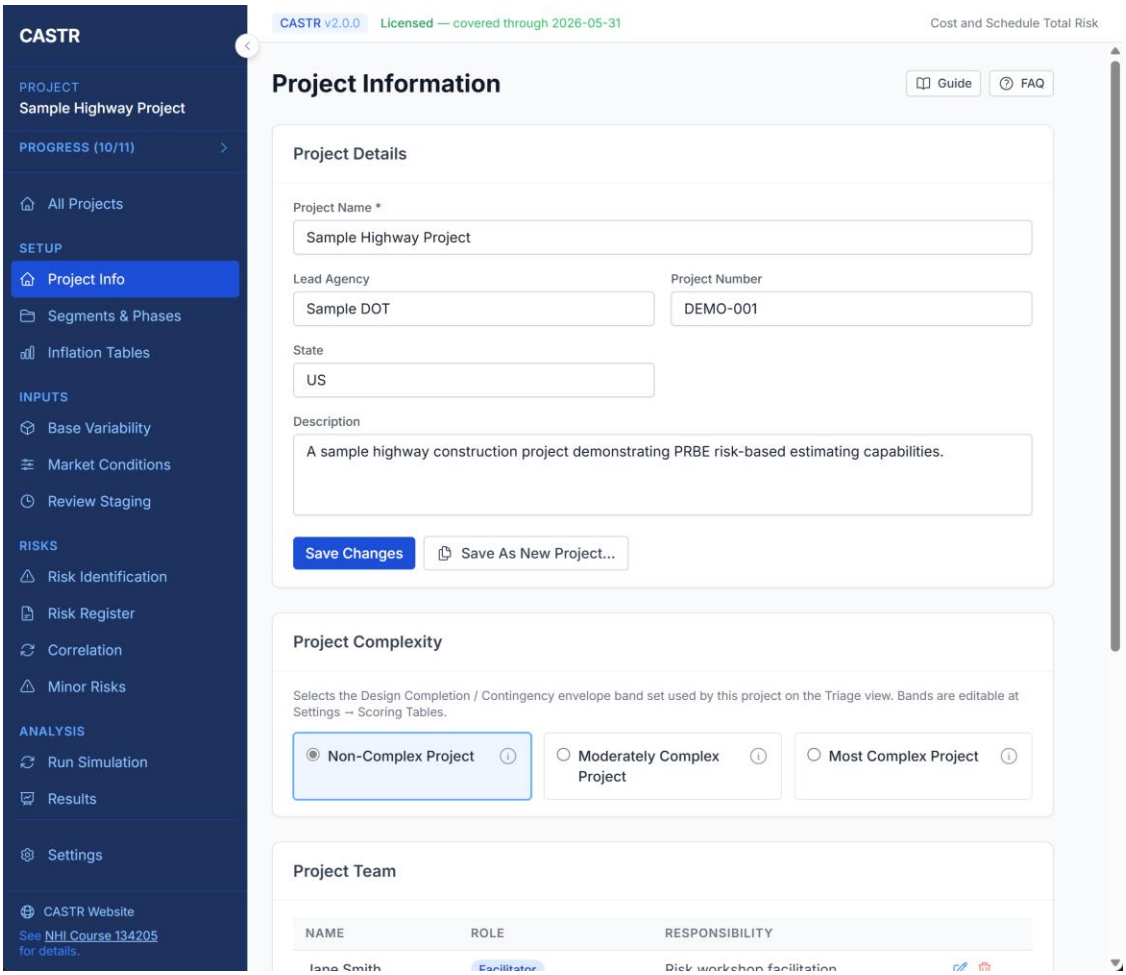

Figure 3.1 - Project Information page

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.

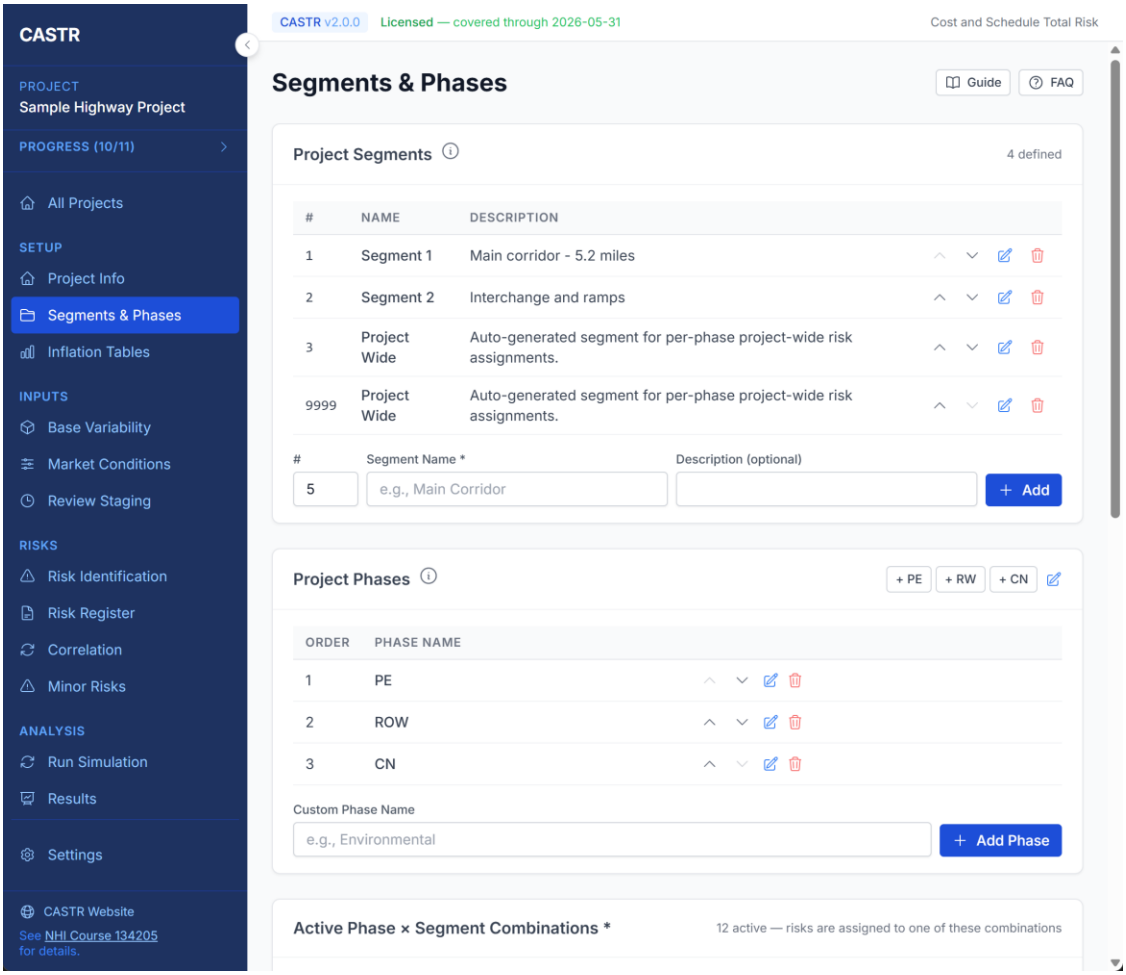

Figure 3.2 - Segments and Phases page

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.

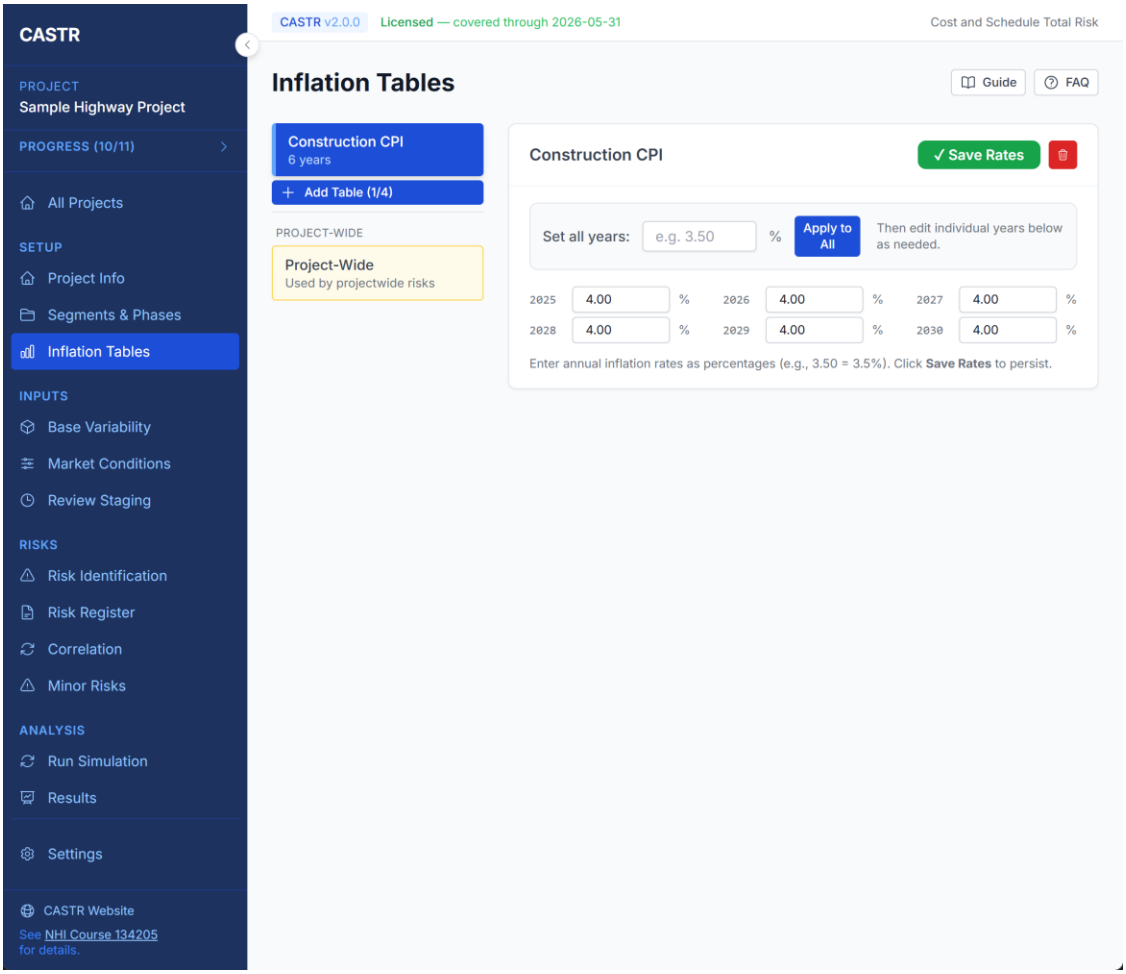

Figure 3.3 - Inflation Tables page

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.

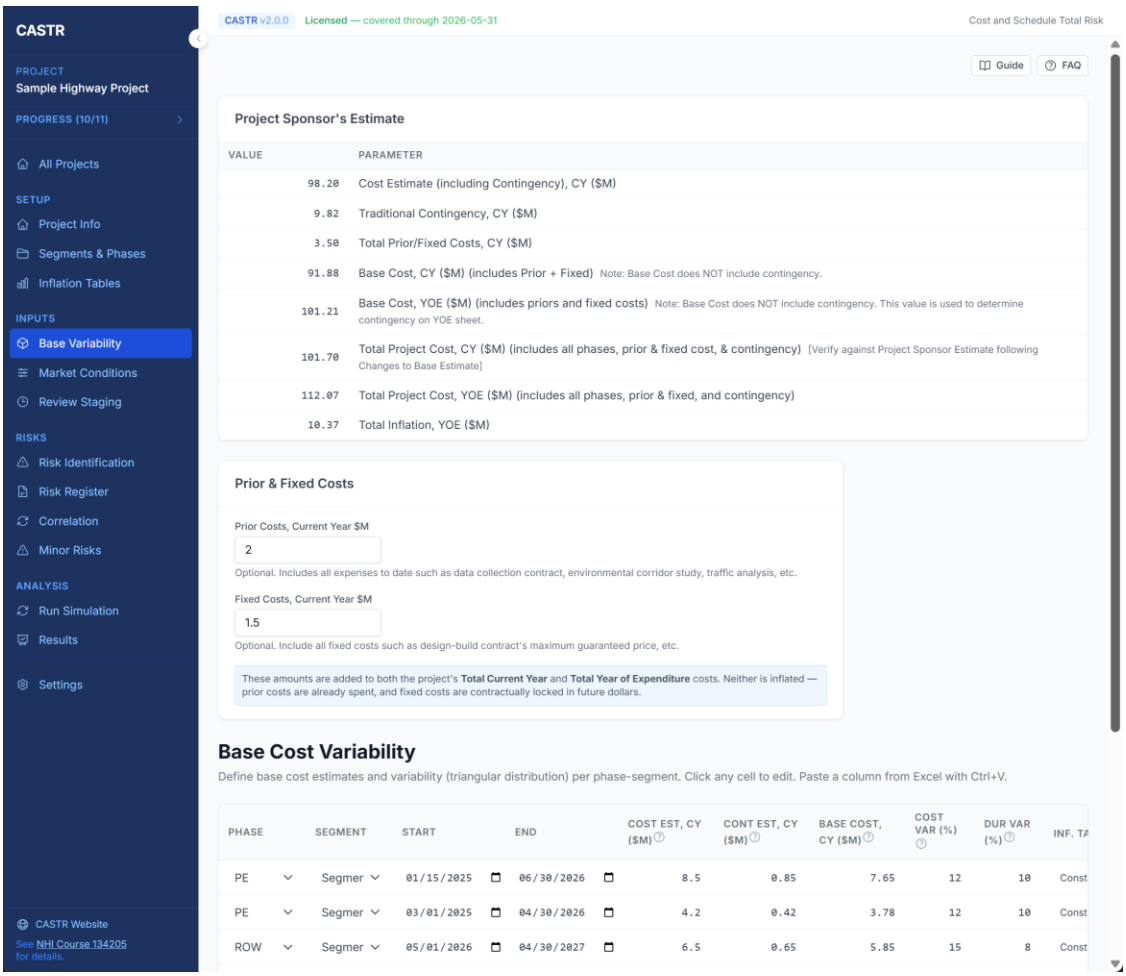

Figure 4.1 - Base Variability page

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.*

4.1.1 Project Wide rows

When a project has **two or more user segments**, CASTR automatically maintains a per-phase **Project Wide** row in the Base Variability table — for example *PE / Project Wide*, *ROW / Project Wide*, *CN / Project Wide* — sitting alongside the per-segment rows for the same phase. These rows are the destination for risks and costs that affect the whole phase rather than one specific segment. When only one user segment exists the Project Wide rows are hidden (a single segment IS the whole phase, so the distinction would be redundant); they re-appear automatically the moment a second user segment is added. Their underlying data is preserved on the backend across that visibility flip, so adding and removing a second segment never loses Project Wide entries you previously made.

Each per-phase Project Wide row has its own **Start Date**, **End Date**, **Cost Est**, and **Inflation Table** like any other row. The shield-S **Snap** button at the right end of each Project Wide row writes the Start and End dates to span all sibling rows in the same phase — earliest start to latest end. See *§16 Tips, Tricks and Best Practices* for the full Snap workflow.

The **Inflation Table** column on a Project Wide row is treated the same way as on a specific phase-segment row: the engine uses whichever inflation table you assigned to that row when computing year-of-expenditure for both the row's own base cost and for any risks targeted at *Phase / Project Wide*. When the Inflation Table cell is left blank, CASTR falls back to the project-wide inflation table designated on the Inflation Tables page (see §3.3). This means base-cost inflation and risk-cost inflation for a Project Wide assignment use

the same table by default, keeping their midpoints aligned — and overriding the table on the row also overrides it for the risks that target it.

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. 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 bundled demo project 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.

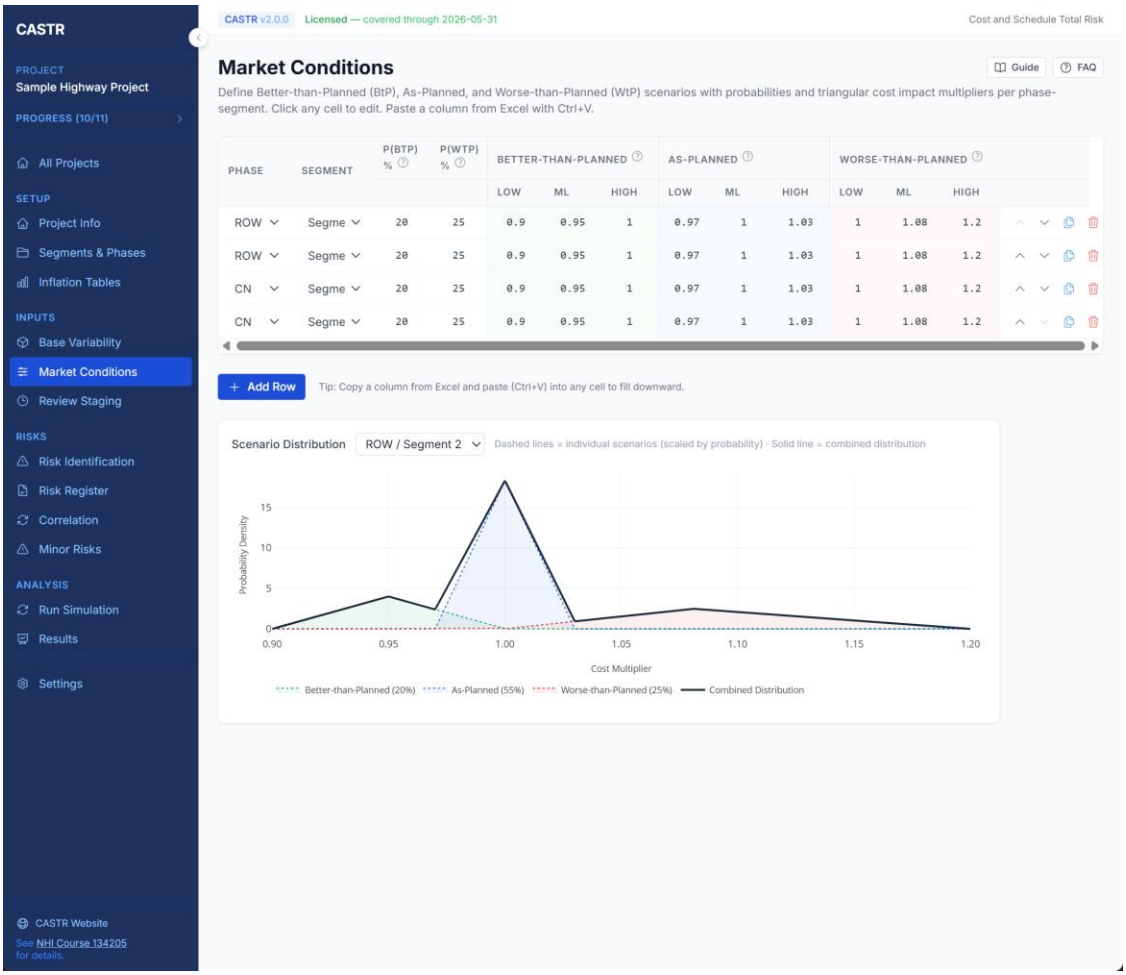

Figure 4.2 - Market Conditions page

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

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

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.

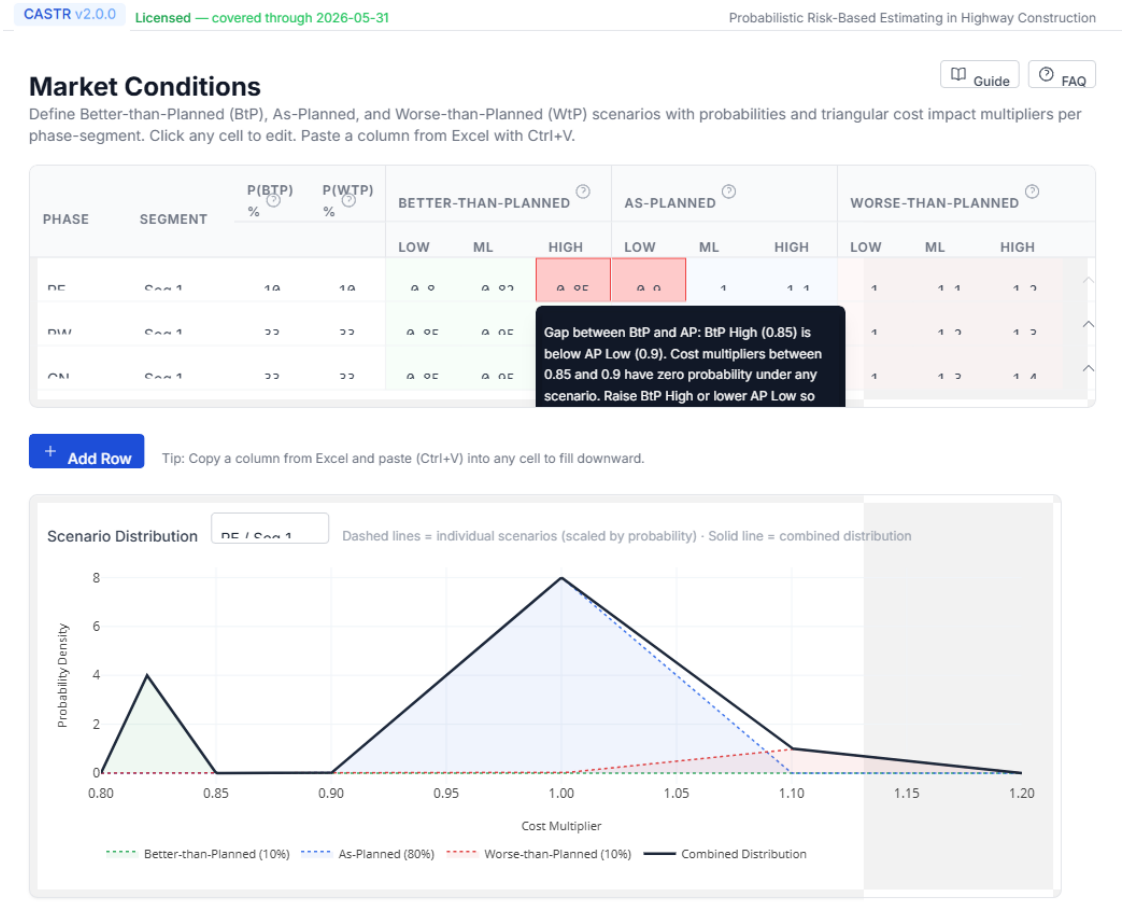

Figure 4.2a - Gap detection: BTP High (0.85) below AP Low (0.90). Both cells are flagged red; the combined distribution in the chart below shows a zero-density region between the BTP and AP triangles.

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

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.

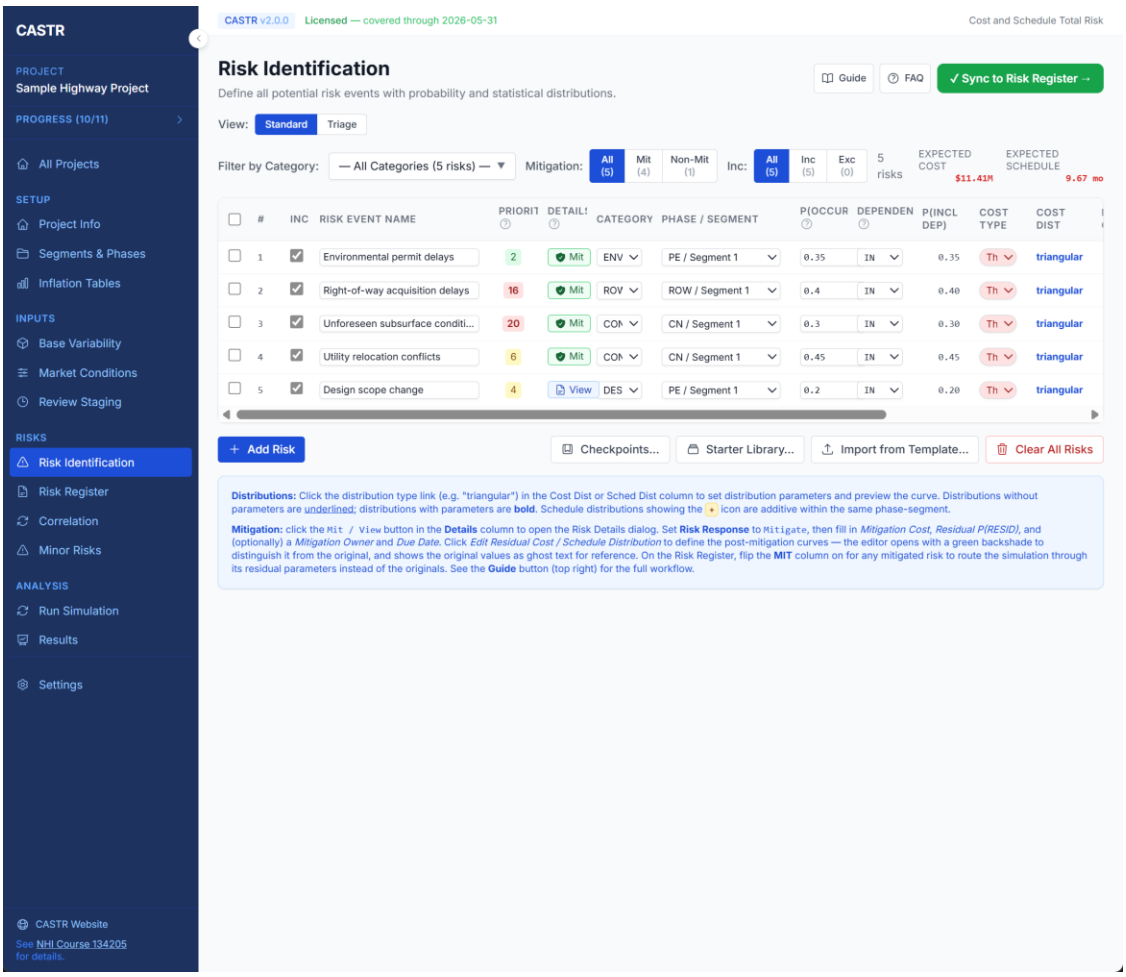

Figure 5.1 - Risk Identification page

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, response configuration, and any work that requires access to documentation fields.

Priority column. 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.

Inc column. The Inc checkbox is the master switch for whether a risk participates in the simulation. Checked (default) — the risk is sampled every iteration and contributes to project totals; the row is eligible for Sync to Risk Register, counted in Priority / Expected-Value totals, and shown in the Results contribution charts. Unchecked — the row stays in the table for documentation but is ignored by the simulator, kept out of the Risk Register, excluded from priority/EV totals, and not shown in contribution charts. Use it to "park" a risk you are still reviewing without deleting it, or to compare results with vs. without a specific risk. Toggling Inc does not save a checkpoint — it is a transient analysis choice.

Help icons. Every column header on both the Standard and Triage views carries a small question-mark icon. Click it to open a popover explaining what the column drives downstream (in the simulation, the Risk Register, or the Excel export). The same convention is used on the Risk Register and the Starter Library admin page.

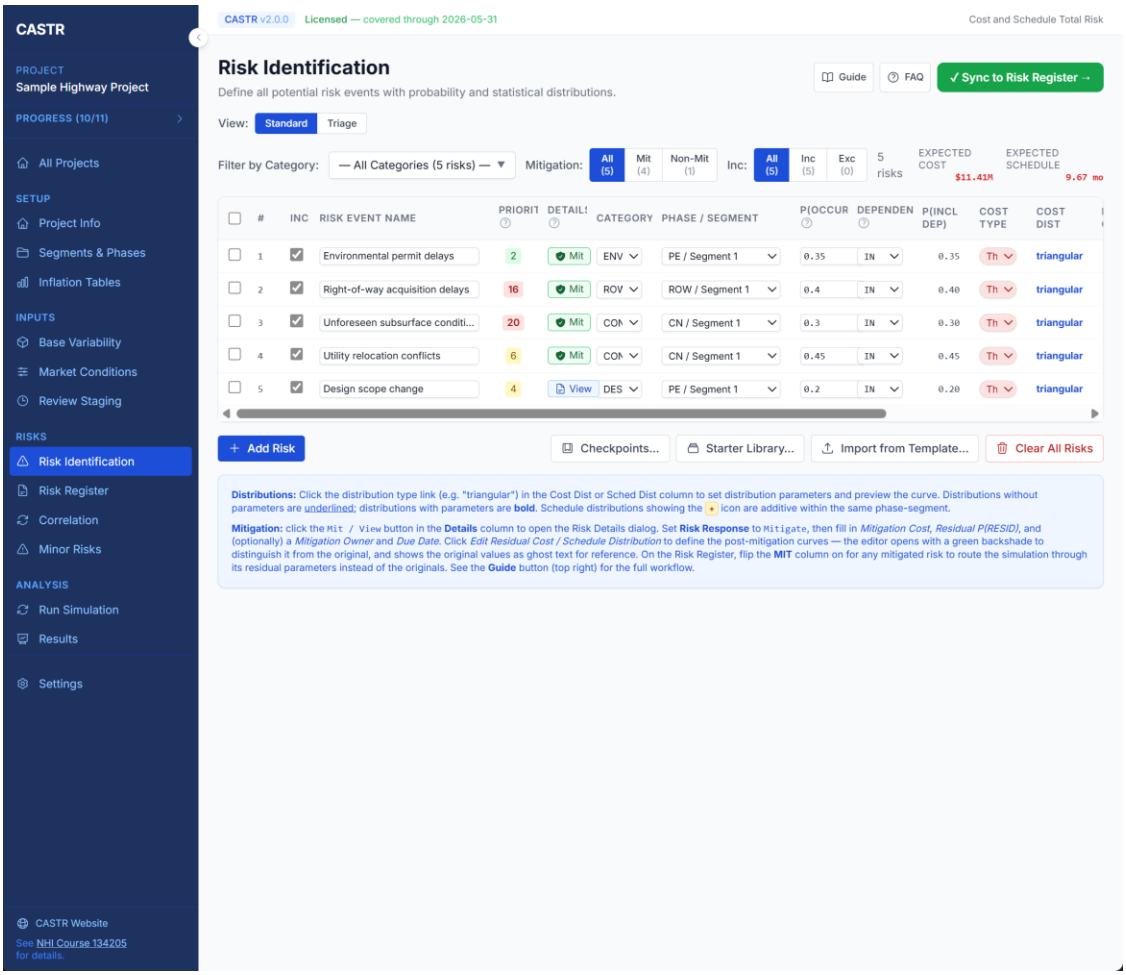

Figure 5.1a - Risk Identification (Standard view) with Priority column

5.1.2 Triage View

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: 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

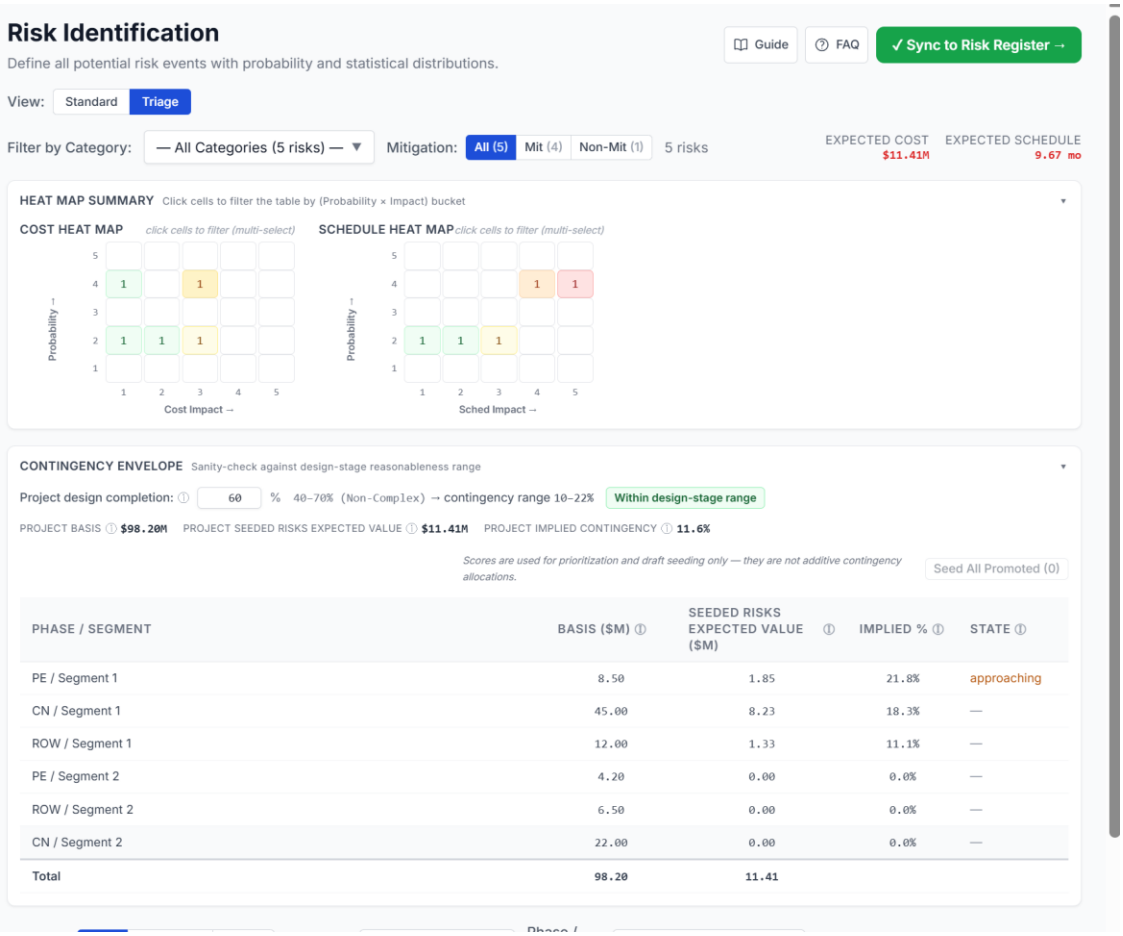

Figure 5.1b - Cost and Schedule Heat Maps (Triage View)

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, response-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

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. 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

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 live on the Standard view envelope card (§5.1.6) and the Risk Register envelope card (§5.2). 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: 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)

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 $(M), SME - ReviewedRisksExpectedValue$(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.

5.1.7 Per-Axis Risk Response

A single risk can carry independent PMI strategies on its cost and schedule axes. The classic case: a market-conditions risk where you plan to MITIGATE the cost-side threat (for example by locking a contract price early) while choosing to ENHANCE the schedule-side opportunity (a faster bid window if the market softens).

Two dropdowns when both axes are populated. When a risk has both a cost distribution and a schedule distribution defined, the Risk Details modal shows two Response dropdowns side-by-side, labeled Cost Response and Schedule Response. When only one axis is populated, only one dropdown appears and it is labeled simply Response, matching the legacy single-axis behavior.

Polarity-appropriate option sets. Each axis dropdown offers the option set that matches its polarity. Threat axes (cost or schedule overrun) offer Accept, Avoid, Escalate, Mitigate,

and Transfer. Opportunity axes (cost underrun or schedule pull-in) offer Accept, Enhance, Escalate, Exploit, and Share. Mixed-polarity risks (threat on one axis, opportunity on the other) draw from the appropriate pool on each side.

Residual flow unlocks per axis. Only Mitigate (on a threat axis) or Enhance (on an opportunity axis) unlocks the residual-distribution flow on that axis — i.e. the Edit Residual Cost Distribution / Edit Residual Schedule Distribution editor and the residual probability field. The other PMI values do not produce a residual model and are not eligible for the Response Analysis pre-vs-post comparison on that axis.

Risk Register Response column. When the two axis values differ, the Response column on the Risk Register shows both, separated by a slash, in the order Cost / Schedule (for example, “Mitigate / Enhance”). When both axes carry the same value the column shows it once (“Mitigate”). When only one axis is populated the column shows that single value.

Resizable Risk Details popup. The Risk Details dialog can be resized by dragging the bottom-right corner (default 672 x ~700 px, minimum 420 x 320, maximum 95vw x 90vh). The header bar remains the drag handle for moving the dialog. Closing on backdrop click is unchanged but no longer fires accidentally if the cursor leaves the dialog during a resize drag.

Notes field. The Risk Details dialog includes an optional free-text Notes field for context that does not fit the structured fields — assumptions you want to record, links to source documents, workshop discussion, etc. The field is hidden by default to keep the dialog clean. Tick the **Add notes** checkbox below the Risk Trigger textarea to expose it; a multi-line textarea appears. Unchecking the box hides the field and clears its content, so hidden text never round-trips silently through saves or exports. When you reopen the dialog on a risk that already has notes, the field auto-expands so the data is always visible. Notes round-trip through the Excel import / export template (Details: Notes column), the Risk Register Notes column, and all register exports (CSV / Excel / Word / PDF). They are also included in Risk Identification snapshots.

5.1.8 Additive vs. parallel schedule risks

The **Additive** column on the Risk Identification table is a per-risk checkbox that controls **how a schedule risk’s delay combines with other schedule risks targeting the same phase-segment** during the Monte Carlo run. It applies only to schedule impact, not cost (cost impacts are always summed across risks). The flag is a property of the risk itself, not of either distribution, and it carries through unchanged into both Pre-Response and Post-Response runs.

The aggregation rule. For each phase-segment, on every iteration of the simulation, the engine sorts that phase-segment’s schedule risks into two buckets based on their Additive flag. It then computes the phase-segment’s total schedule delay for that iteration as:

delay = max(sum of additive-bucket delays , max of parallel-bucket delays)

In words: additive risks stack with each other, parallel risks compete with each other (only the worst applies), and the phase-segment ends up taking whichever of those two combined values is larger.

Why parallel is the default. In construction-management practice, two activities running concurrently inside the same phase-segment do not *both* extend calendar time when one is delayed — the contractor shifts crews, equipment, and shifts to keep the other moving. Summing every concurrent delay would systematically overstate total schedule risk. CASTR therefore treats schedule risks on the same phase-segment as competing for a single calendar window (max), and you opt in to summation only when the activities truly chain end-to-end on the critical path. New risks default to **Additive = off** for this reason.

Worked example. Three schedule risks all target `CN / Seg 1`. On a single Monte Carlo iteration where all three fire, their sampled delays are 2, 10, and 6 months. The phase-segment's contribution to project duration on that iteration depends on the flag:

Configuration	Additive sum	Parallel max	`max(...)` = phase-segment delay
All three Additive = off (default)	—	max(2, 10, 6) = **10**	**10 months**
All three Additive = on	2 + 10 + 6 = **18**	—	**18 months**
The 10-month risk Additive = on; the 2- and 6-month risks Additive = off	10	max(2, 6) = 6	max(10, 6) = **10 months**

The fully-parallel case gives 10 months because the contractor would shift resources around the smaller delays so they overlap with the worst one. The fully-additive case gives 18 months because the activities are assumed to be sequential — each adds its own calendar time. The mixed case shows how the rule reads when both buckets are populated.

When to flip Additive on. Use it when the risks describe events that each independently consume their own block of calendar time on the critical path:

- Sequential review or approval steps that must each complete before the next starts (for example, *agency review delay* followed by *FHWA concurrence delay* — both required, neither parallelizable).
- Distinct rework cycles that must follow one another rather than overlap.
- Successive weather windows or shutdown periods that aren't substitutable for each other.

When to leave Additive off. Most schedule risks are concurrent — they describe different aspects of the same calendar window inside one phase-segment, where a contractor can

absorb one delay by shifting attention to another front. Leave Additive off whenever there is realistic float to do so.

Edge case — project-wide risks. A risk with no phase-segment assigned (target_sp_id = none, also called "Project Wide") ignores the Additive flag and is **always summed** with the other project-wide risks. Project-wide risks are assumed to affect the whole-of-project schedule and are not eligible for the parallel/max rule, since there is no shared phase-segment calendar window for them to compete over.

How to read this on Results. The Tornado chart and Tail Drivers ranking show each schedule risk's individual variance contribution; they do not visually distinguish additive from parallel risks. To check the impact of the flag on a project, toggle it on the suspect risk in the Risk Register and re-run. The Compare-Both run mode is handy here — pre/post curves let you read the Total Project Duration delta directly.

5.1.9 Bulk operations — Clone Selected and Delete Selected

The SEL column on the leftmost edge of the Risk Identification table holds a checkbox on every selectable row. Tick any number of rows to select them; the header checkbox toggles every visible, unlocked row at once. Rows locked in the Risk Register are not selectable (the bulk endpoints reject them and the checkbox is suppressed).

When the selection is non-empty, two buttons appear next to Add Risk below the table:

- **Clone Selected (N)** — copies every selected risk to a new row at the bottom of the table. Each clone is a complete copy of its source: probability, distribution parameters, category, phase-segment, additive flag, dependency, residual values, and all Details fields (including Notes) are preserved. Only the row number is re-assigned and the Risk Register sync number is cleared (so the clones can be synced as fresh register rows). After cloning, the new rows are automatically selected so you can immediately edit them — typically to set a different phase-segment on each clone to model "one underlying risk that affects multiple phases". A checkpoint is auto-saved before the operation and the new clones are *not* automatically correlated; if the clones should fire together in the Monte Carlo, link them at coefficient = 1.0 in the Correlation Editor (Cost ↔ Cost and Sched ↔ Sched).
- **Delete Selected (N)** — removes every selected risk in a single transaction. A checkpoint is auto-saved first so the action is reversible from the Checkpoints dialog. Rows locked in the Risk Register cannot be deleted; the endpoint refuses the operation rather than partially completing it.

Both buttons prompt for confirmation before acting. The Risk Identification table preserves your selection through edits and search, so workshops can stage a selection across multiple filter passes and act on it once.

5.2 Risk Register

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.

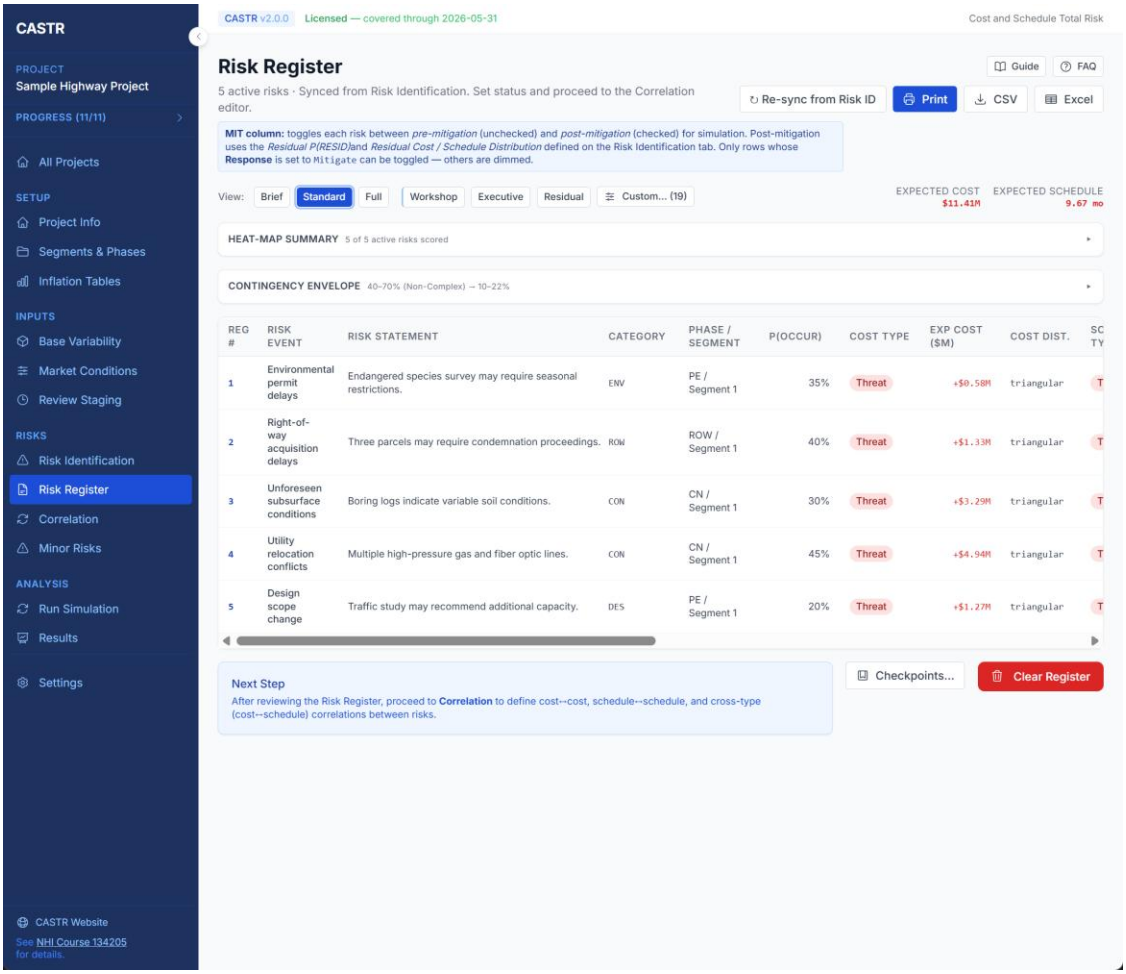

Figure 5.2 - Risk Register page

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 (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), $SME-ReviewedRisksExpectedValue$(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.

RES auto-baseline. When you click Sync to Risk Register the system automatically sets RES = checked on any new row whose linked Risk Identification event has Risk Response = Mitigate OR has any residual value defined (residual probability, residual cost distribution, or residual schedule distribution). Existing rows are also auto-promoted from unchecked to checked at sync time when their residuals are present. The premise: if an SME has bothered to define a residual, the working assumption is that the response should be in the analysis. The RES checkbox is editable whenever Response = Mitigate OR any residual is defined; it is dimmed otherwise (no response has been modeled, so the toggle would have no effect).

Untoggling persists. Once you uncheck a row, that choice persists across page reloads, refetches, and other passive reads. The only way for RES to flip back to checked is an explicit Sync to Risk Register click while the linked event still has residuals defined. If you want a residual-priced risk to stay RES-OFF after Sync, just leave it that way.

Response column visibility. The Response column group (RES, Response, P(RESID), residual distributions, etc.) is auto-hidden when no risk on the register has Response = Mitigate AND no risk has residual values defined. As soon as either condition is true on at least one row, the columns appear automatically.

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.

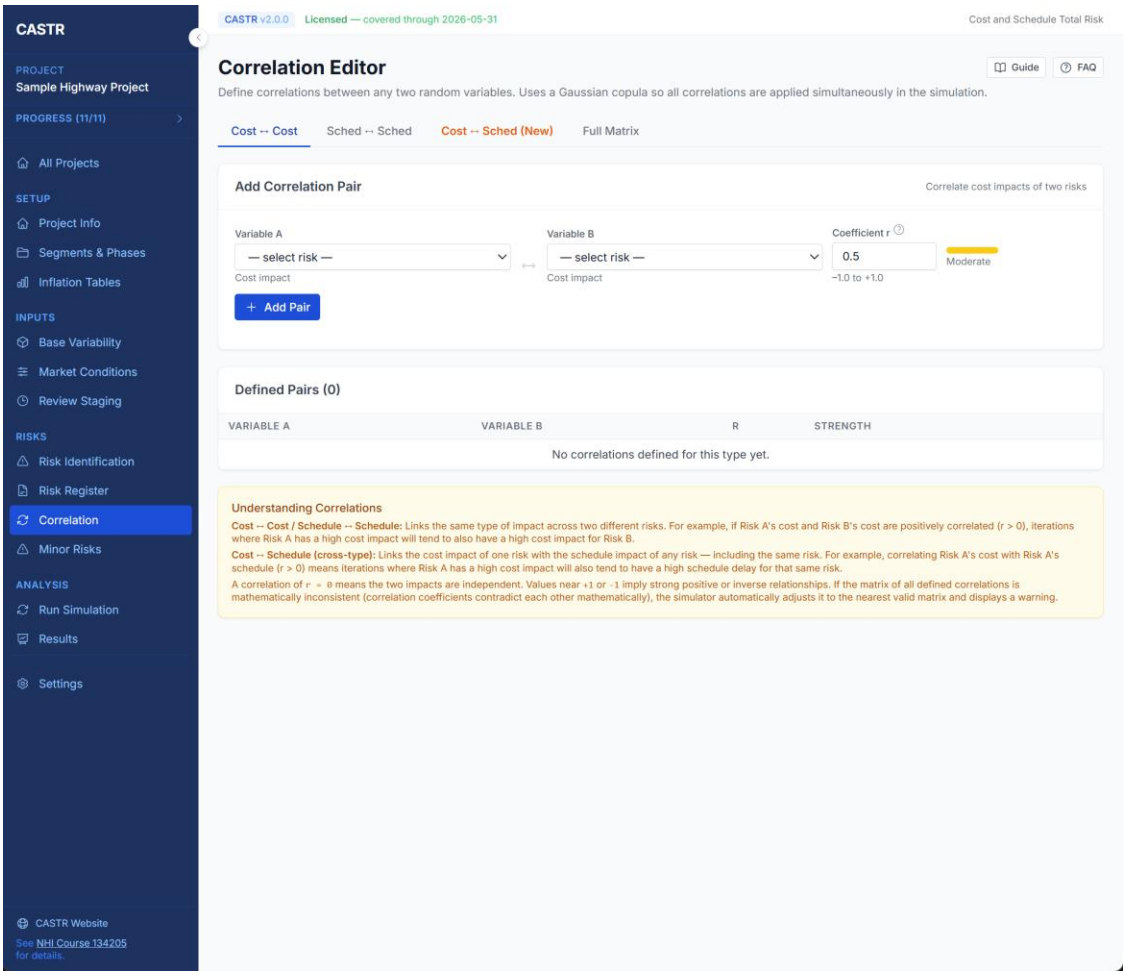

Figure 5.3 - Correlation Editor page

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), 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. 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.

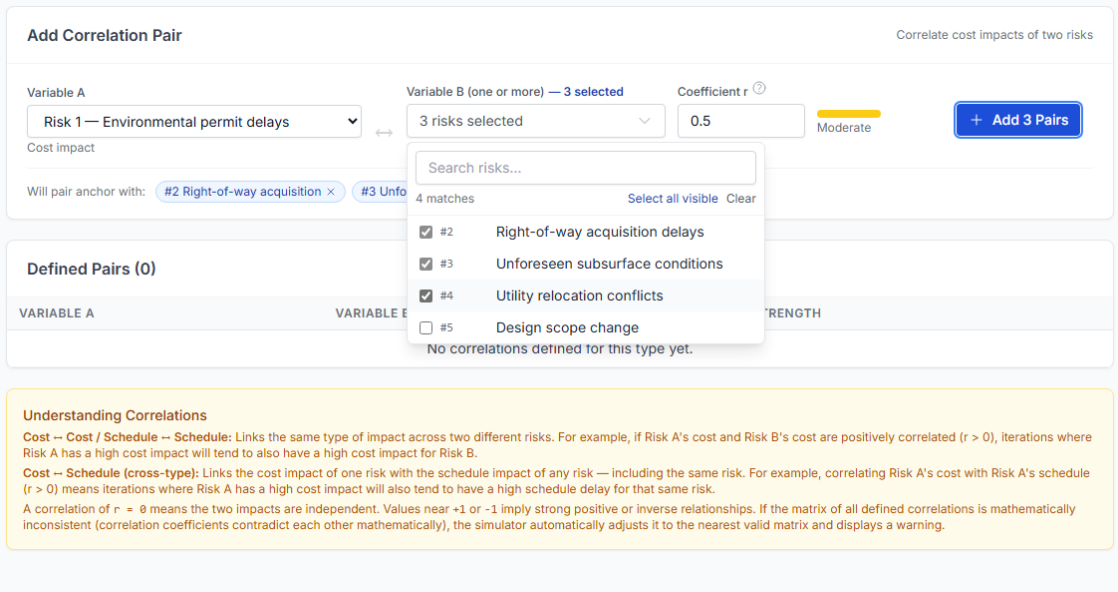

Figure 5.3a — One-to-Many correlation entry. Variable A is a single anchor; Variable B is a multi-select picker. The chip row below the form previews exactly which (anchor, target) pairs will be written on Save. The Add button label reflects the count (here: "Add 3 Pairs").

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.

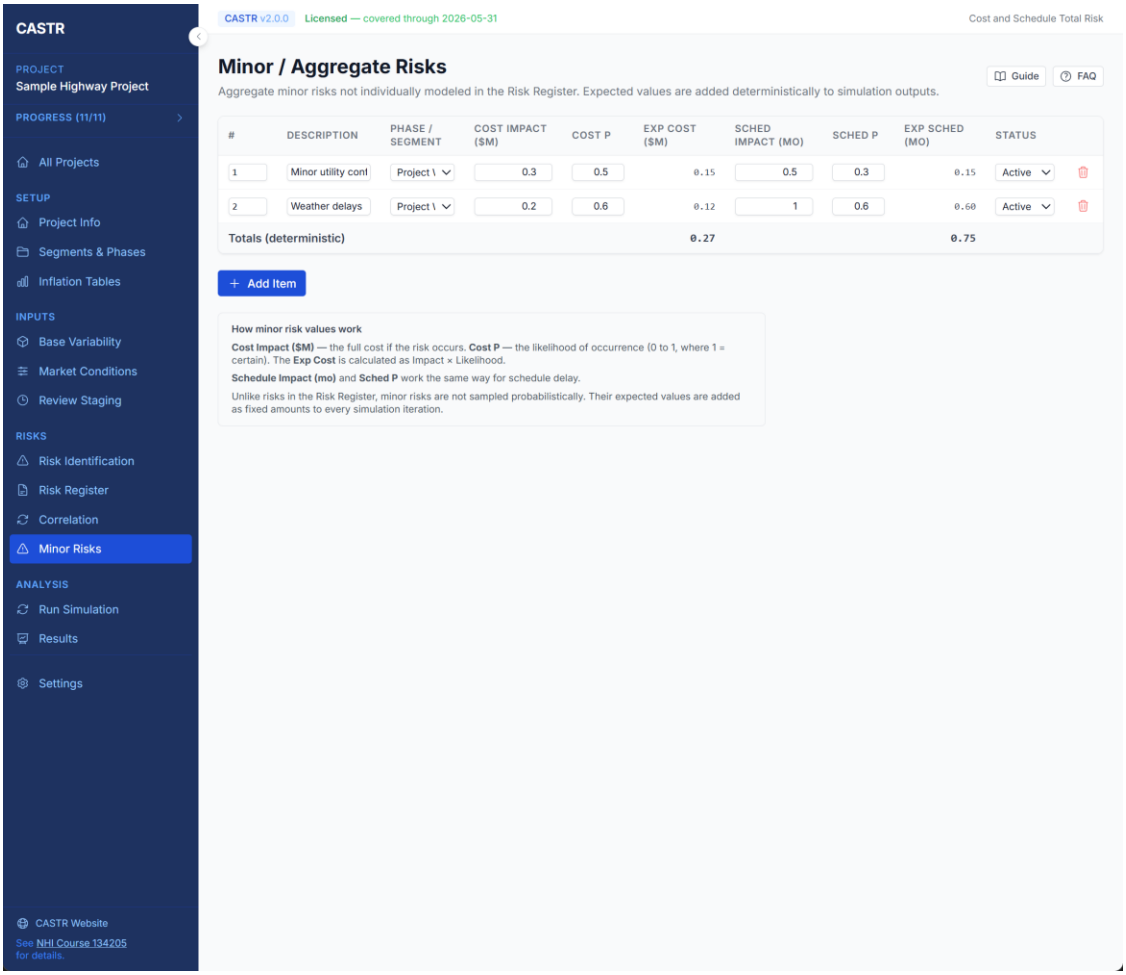

Figure 5.4 - Minor / Aggregate Risks page

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.

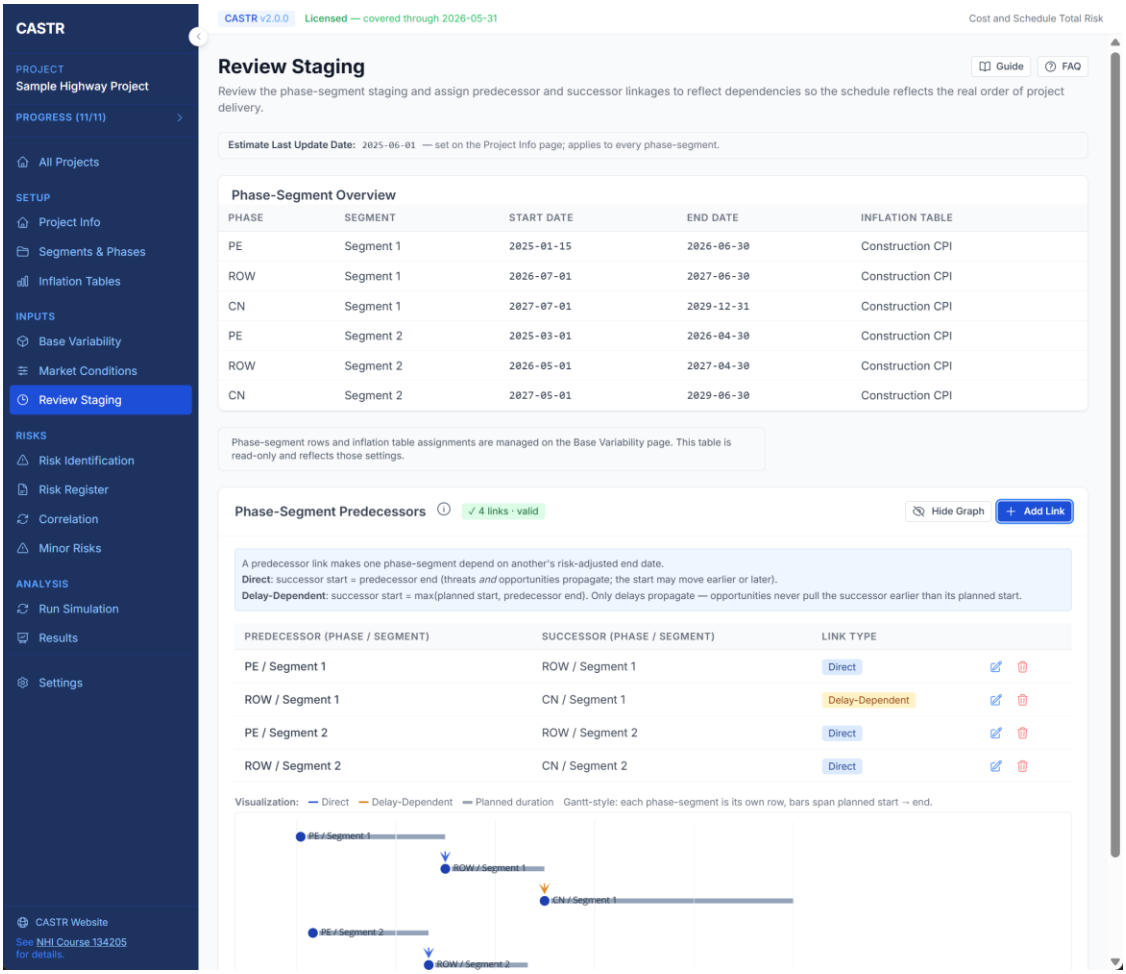

Figure 5.5 - Review Staging page

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.

Why predecessor links matter for response analysis. Total Project Completion Date and Total Project Duration are both built by walking the predecessor graph defined in this section. If your project has parallel phase-segments (for example, PE running alongside ROW with both feeding into CN) and you do NOT define predecessor links, the simulator treats every row's start date as fixed at its planned start. A delay to PE then has no effect on CN's start, and the project completion date stays pinned to whichever row's planned end is largest. The result: a response that targets a non-critical-path risk reduces the row's own duration but leaves the Completion Date unchanged. If your pre vs post response curves move on Duration but not Completion Date, the most likely cause is missing or incorrect predecessor links — wire them here using the Direct (D) link type and re-run the simulation.

Chapter 6: Running the Simulation

6.1 Simulation Settings

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.

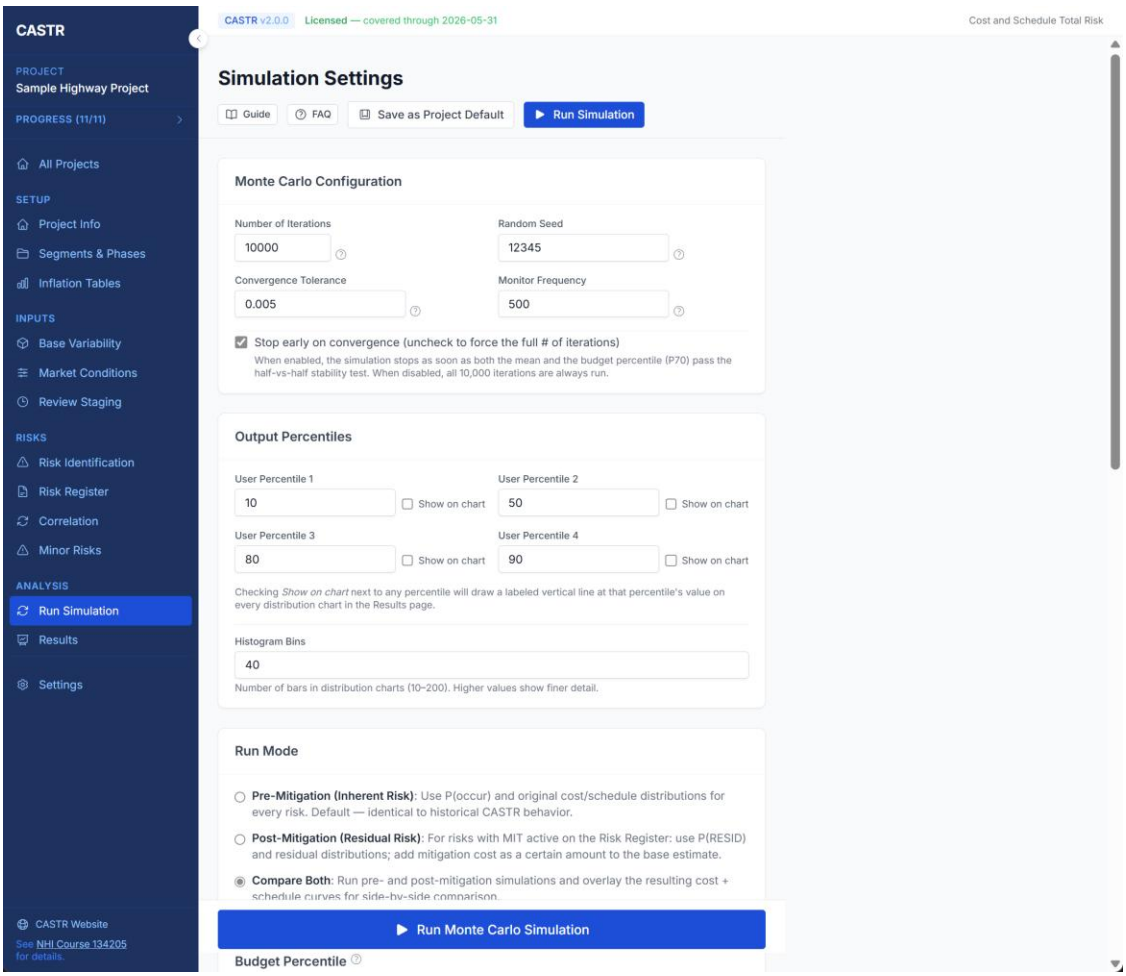

Figure 6.1 - Simulation Settings page

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. 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).

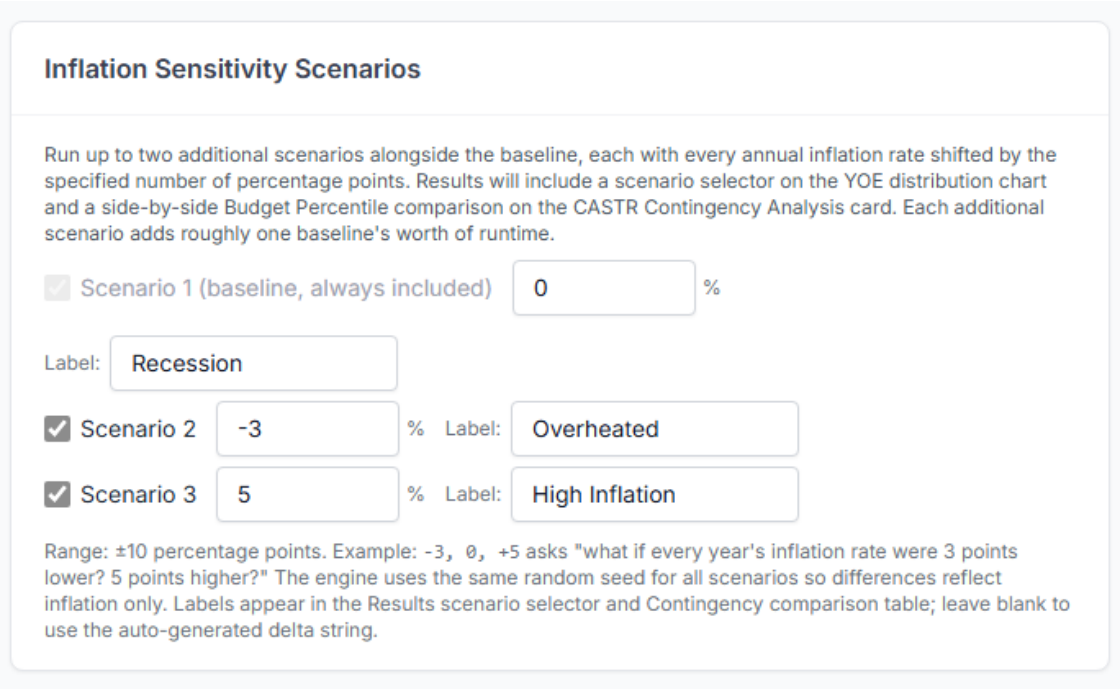

Figure 6.2 — Inflation Sensitivity Scenarios panel on the Simulation page. Scenario 1 is the always-on baseline (delta = 0). Up to two additional scenarios can be enabled with their own deltas (in percentage points) and optional labels. Empty labels fall back to the auto-generated delta string (e.g.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

6.4 Early-Stop Convergence and Tail Drivers

When the “Stop early on convergence” checkbox is on, the simulator monitors the YOE total cost stream every Monitor Frequency iterations and stops the run when both the mean and the budget percentile are stable in a half-vs-half ratio test. This is appropriate for the headline contingency number but is, on its own, a weaker guarantee for the per-risk rankings on the Tail Driver Analysis card.

Tail-aware floors.

The convergence test will not declare convergence until each half of the iterations has at least 200 samples in the budget-percentile bucket AND at least 200 samples in each tail-driver bucket (the percentiles set under “Tail Driver Analysis” — defaults P10 and P90). This prevents the run from passing the ratio test by sheer noise on a thinly-sampled tail. Required iterations vary by setting:

Default settings (P70 budget, P10/P90 tails) — floor 2,000 iterations; runs typically converge somewhere in the 6,000–10,000 range.

Budget P95 — floor 8,000 iterations.

Budget P99 — floor 40,000 iterations.

Tail percentiles P5/P95 — floor 4,000 iterations.

Tail percentiles P2.5/P97.5 — floor 8,000 iterations.

Per-risk fire counts in the tail.

The convergence test settles the OVERALL tail bucket. Each individual risk only fires inside that bucket at its own probability rate. With default settings, a 30%-probability risk fires ~90 times in a 300-sample tail (rankings reliable), but a 5% risk fires only ~15 times (rankings noisy). If the headline driver of a tail you care about is a low-probability risk, the early-stopped run may not have enough fires for that risk's ranking to be statistically stable.

The Results page Tail Driver Analysis card surfaces this directly via a Reliable / Review / Insufficient confidence chip on every tail block (see §7.4). When the chip on a tail you care about reads Review or Insufficient, untick "Stop early on convergence" and re-run with a higher iteration count; rough rule of thumb, a top-3 driver needs ~100 fires inside the tail bucket to settle, so total iterations needed are roughly 100 / (P_occur × tail_width). A 5%-probability risk in a 10% tail bucket needs ~20,000 iterations.

Recommended early-stop policy.

Workshop runs / quick draft: leave early-stop ON. The default settings produce defensible headline numbers and good rankings for high-probability tail drivers.

Reporting / formal contingency memos: turn early-stop OFF and run the full iteration count. Reproducibility is more important than wall time at this stage.

Projects with low-probability / high-impact risks in the tail ("tail-busters"): turn early-stop OFF, raise the iteration count to 20,000–50,000 if the confidence chip on the tail you care about reads Review or Insufficient at the default count.

Budget percentile set above P90: leave early-stop ON — the new tail-aware floor will refuse to stop early until the budget percentile itself is well-sampled, which automatically provides ~the same protection.

Chapter 7: Understanding Results

7.1 Summary Statistics

The Results Summary tab provides a high-level overview of the simulation outcomes — the recommended risk-based contingency at the top, the Distribution Health diagnostics card next, the Inflation Sensitivity comparison card (when multi-scenario), then the **Summary Statistics combined table**, and the per-phase-segment cost statistics cards at the bottom. The Joint Confidence chart lives on its own dedicated tab — see §7.8.

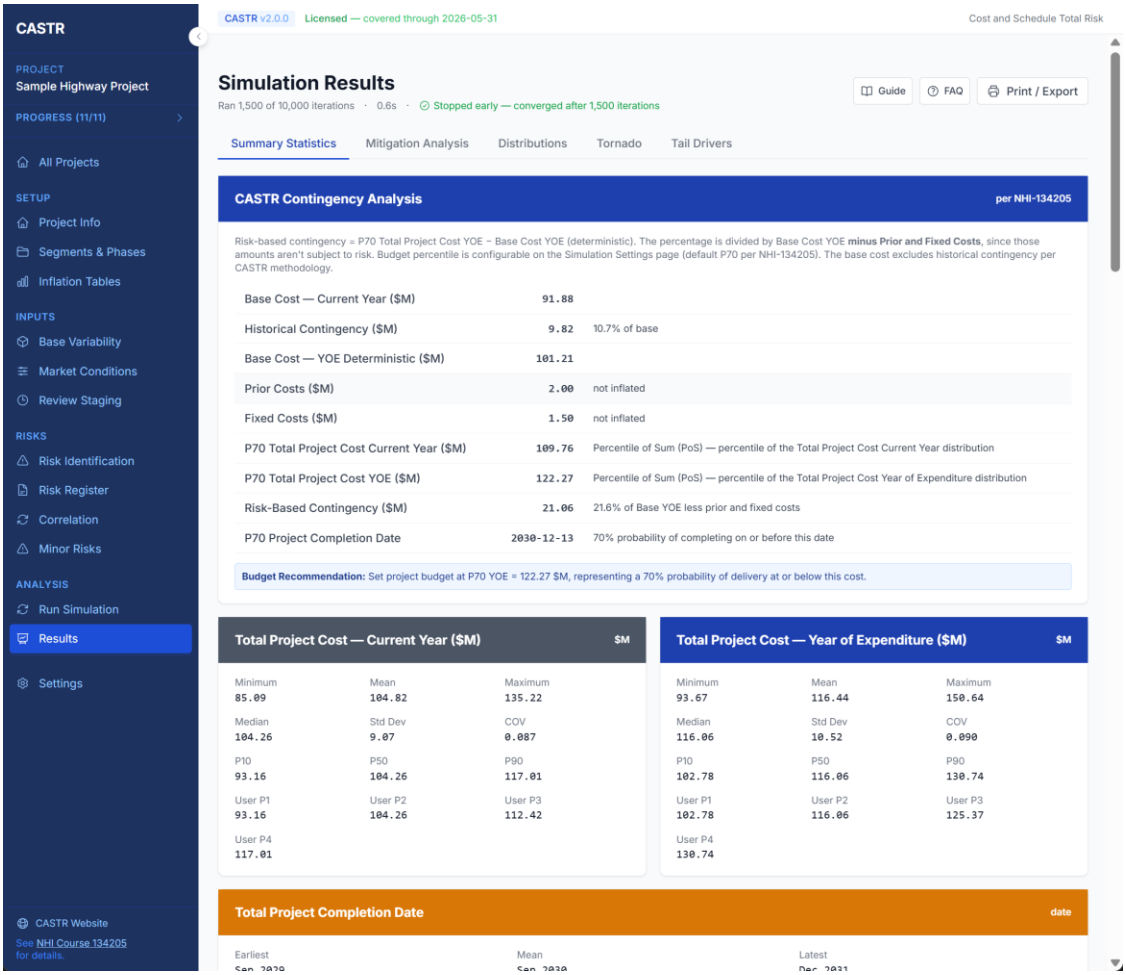

Figure 7.1 - Results Summary Statistics

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.

Contribution to Budget waterfall. The Contribution to Budget card directly below the Contingency Analysis card walks the same budget-percentile Total Project Cost YOE step-by-step from the deterministic current-year base up to the final budget figure. It is the picture that answers the question every executive sponsor eventually asks: *"what is my budget actually made of, and how did we get to that number?"*

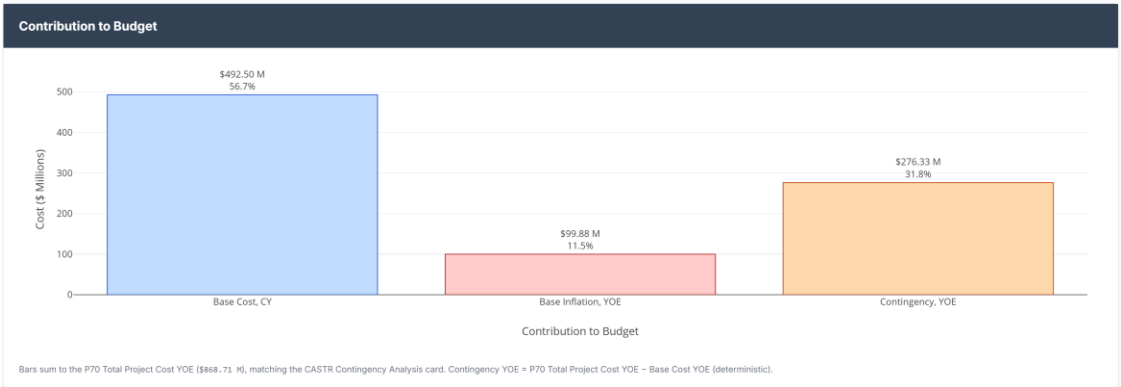

Figure 7.1a — Contribution to Budget. A waterfall that starts at Base Cost CY (light blue, drawn from zero), adds Base Inflation YOE (light red), then adds the contingency components (light orange) — either a single Contingency YOE step or, when Component Analysis has been run, three sub-steps for Base Variability, Market Conditions, and Risk Register — and finishes at the budget-percentile Total YOE bar (light blue, drawn from zero), which equals the headline budget shown on the CASTR Contingency Analysis card above.

The steps in the waterfall are:

- **Base Cost, CY** — the deterministic current-year project cost (includes prior costs and fixed costs). The work itself, in today's dollars, with no contingency and no inflation.
- **+ Base Inflation, YOE** — the deterministic cost lift from inflating each phase-segment to the midpoint year of its scheduled execution. Prior and fixed costs are not inflated, so they cancel between Base Cost YOE and Base Cost CY and contribute zero here.
- **+ Contingency, YOE** — the risk-based contingency, calculated as the budget-percentile Total Project Cost YOE minus the deterministic Base Cost YOE. The reserve being held to deliver at the chosen confidence level. When Component Analysis has been run on the Run Simulation page, this single orange step splits into three sub-steps — **+ Base Variability**, **+ Market Conditions**, **+ Risk Register** — each rescaled to sum exactly to the on-page Risk-Based Contingency value. Without Component Analysis the waterfall collapses to the single aggregate Contingency step using the same total.
- **{Pbudget} Total YOE** — the final bar, drawn from zero, equal to the sum of all preceding steps. This is the budget-percentile Total Project Cost YOE shown on the Contingency Analysis card. Light-blue like Base Cost CY because both are project bookends.

Dotted connectors join the top of each step to the base of the next so the running cumulative total is visible at a glance. The chart is guaranteed to sum to the budget-percentile Total Project Cost YOE on the Contingency Analysis card above; the caption beneath the chart re-states this identity with the actual numbers from the current run.

Color convention. Light blue for the project bookends (Base Cost CY and the final Total YOE bar). Light red for inflation — kept visually distinct from contingency because **inflation is NOT contingency**. Inflation is the deterministic price-level lift between the cost-estimate date and the year of expenditure; it would be there even on a zero-risk project. Light orange for contingency. Contingency exists for RISK only — base variability, market conditions, and risk register events. The color separation prevents the two from getting conflated in a leadership briefing.

Why this matters for leadership briefings. Boards, mayors, and finance committees rarely want to argue about Spearman correlations or P70-versus-P80. They want to know what the request is *for* and how the team got to it. The Contribution to Budget waterfall frames the answer step-by-step — *"this is what the work costs today, this is what inflation will add by the time we spend the money, this is the reserve we are holding to absorb risk we know but cannot pin down, and (with Component Analysis) here is the share of that reserve driven by base estimating variability, by market conditions, and by specific risks on the register."* In that order, those steps are the single most defensible way to walk a non-technical audience through the budget number CASTR is recommending. Pair the chart with the headline percentile from the Contingency Analysis card and most audiences need nothing else to sign off on the request.

Summary Statistics combined table. The Total Cost CY, Total Cost YOE, and Total Project Completion Date statistics are presented in a single, dense table. Each *row* is one statistic; each *column* is one project metric. The row whose percentile matches the agency-configured Budget Percentile (P70 by default) is highlighted in pale yellow.

Rows (24 statistics). Min, Max, Mean, Median, Standard Deviation, Coefficient of Variation, Skewness, Kurtosis, P10–P90 Spread, the four user-defined percentiles (configurable on the Simulation Settings page — typically P10 / P50 / P80 / P90), and the eleven deciles P0 through P100.

Columns (seven metrics). The four "Total Project" columns are the headline outcomes you defend in the workshop:

- **Total Project Cost — CY ($M)** — the simulated total in current-year dollars (no inflation).
- **Total Project Cost — YOE ($M)** — the same total in year-of-expenditure dollars, with each phase-segment's cost inflated to its midpoint date.
- **Total Project Duration (months)** — calendar time = project completion date − earliest planned start.

- **Total Project Completion Date** — the simulated completion-date distribution. (Skewness and kurtosis are computed on the date ordinal samples and so are unitless and directly comparable to the cost columns.)

The three "Total Project Risk Register" columns isolate the contribution of risk-register events alone — they exclude base variability and market conditions. The difference between a "Total Project" column and its corresponding "Risk Register" column tells you how much of the variability comes from non-register sources:

- **Total Project Risk Register Delay (months)** — per-iteration aggregate of schedule risks (project-wide bucket plus per-segment combined deltas, where parallel risks aggregate via max).
- **Total Project Risk Register Cost — CY ($M)** — sum across all risk events of the signed (Threat = +, Opportunity = −) cost impact per iteration.
- **Total Project Risk Register Cost — YOE ($M)** — same, with each risk's impact scaled by its phase-segment's YOE inflation factor.

Note that the values in the Risk Register columns can be derived from the Total Project columns and the deterministic base cost — for example, Risk Register Cost CY = (Total Project Cost CY) − (Base Cost CY) − (base-variability contribution) − (market-condition contribution). Splitting them out explicitly lets you answer "what share of my contingency is the register driving?" without doing the arithmetic in your head.

How to read it. Compare Mean vs. Median (P50) — if Mean is materially higher, you have meaningful tail risk and should defend at P70 or higher rather than P50. The COV row tells you the spread relative to the mean (standard deviation ÷ mean) — treat it as a descriptive statistic only, not as a calibration benchmark. Skewness > 0 means a right tail (typical for cost); kurtosis > 0 means heavier tails than Normal. See §12.5 for full interpretation rules.

Multi-scenario runs. The Total Project Cost YOE column re-renders against the inflation scenario you pick from the chip selector above the table. Cost CY, Duration, and Completion Date are scenario-invariant by design. The Risk-Register YOE column reflects the baseline scenario only; flipping the chip does not change it.

Density. The table uses a deliberately tight row height so all 24 rows fit on screen without scrolling on a typical workshop laptop. The Statistic column is sticky on horizontal scroll so you can read the row labels as you pan across the seven metric columns.

Help icon. The header carries a question-mark icon that opens a popover summary plus a deep link into this manual section.

7.2 Distribution Charts

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

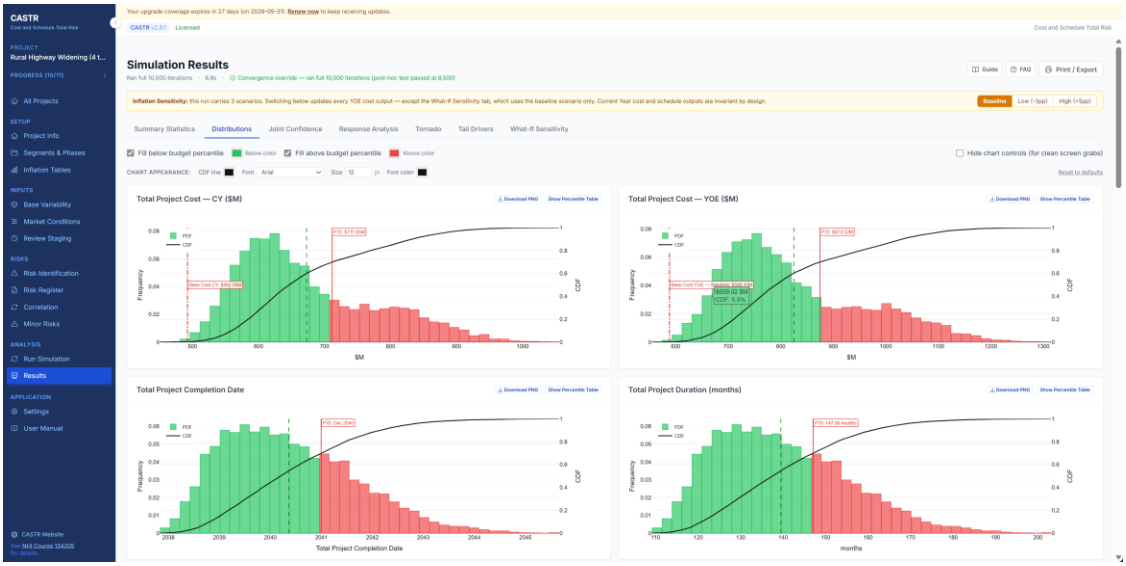

Figure 7.2 - Results Distribution Charts

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. 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 response efforts will have the most impact.

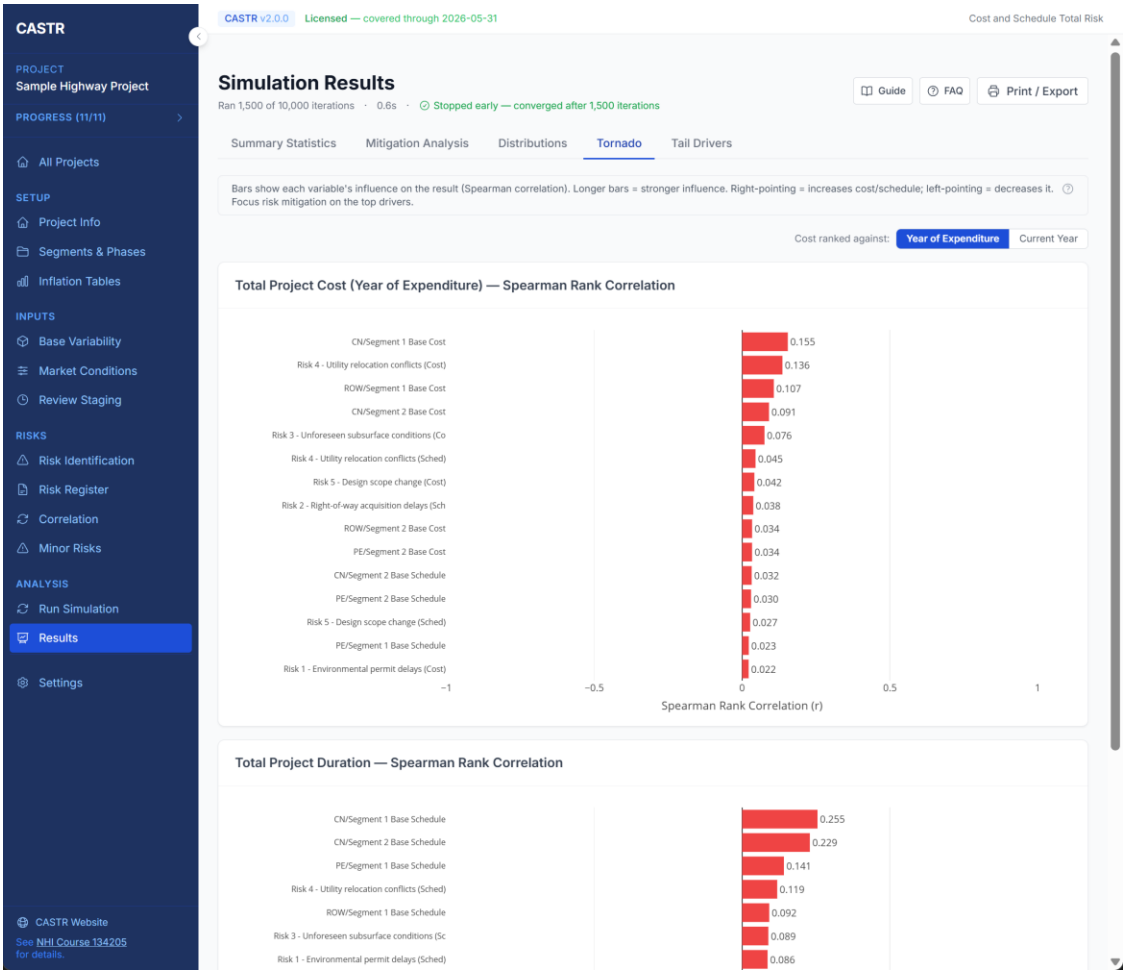

Figure 7.3 - Results Tornado Charts

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. 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.

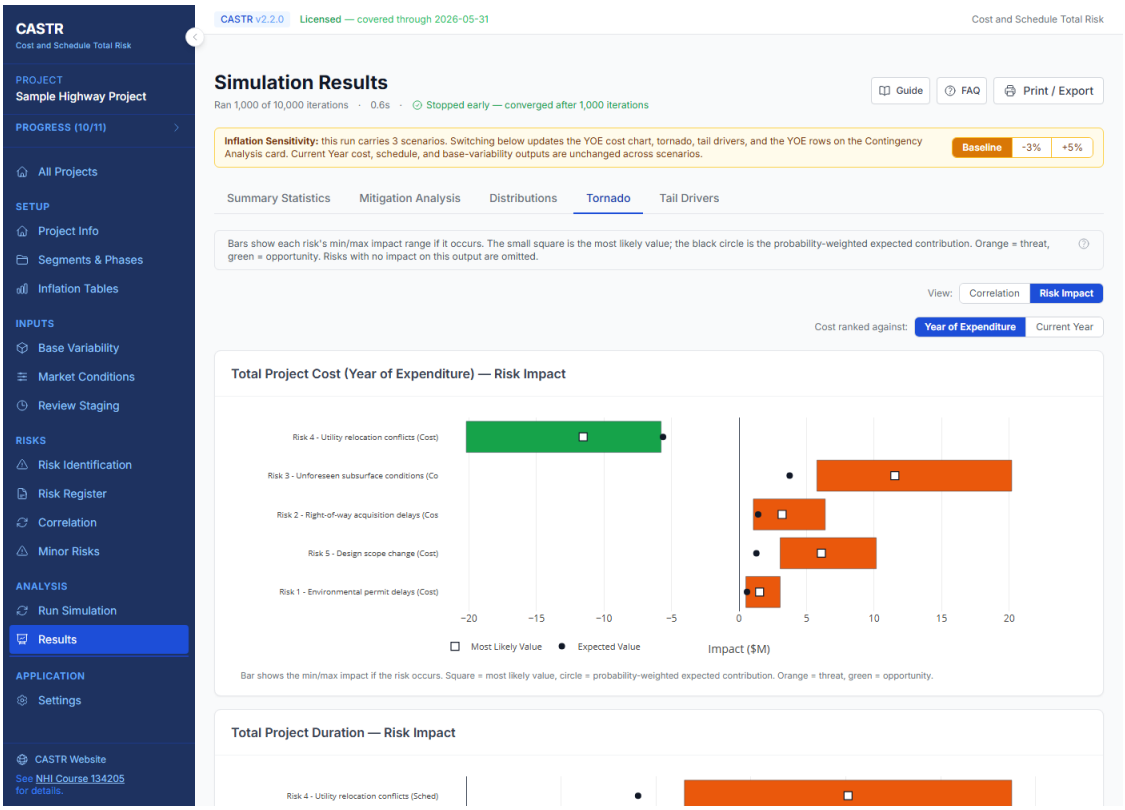

Figure 7.3a — Tornado in Risk Impact mode. Bars are signed expected dollar contributions to mean cost, ranked largest first. Use the toggle in the upper right of the chart to flip between Spearman correlation and Risk Impact.

7.4 Tail Driver Analysis

Tail-driver confidence chips. Each tail card now carries a chip in its header — Reliable (green), Review (amber), or Insufficient (red) — that tells you how trustworthy the per-risk rankings inside that card are. The number in parentheses next to the chip label is the lowest fire count among the top-3 drivers (the bottleneck driver — the one with the fewest data points behind its position).

Reliable means every top-3 driver fired ≥ 100 times in the tail bucket; rankings are stable across re-runs and defensible. Review means at least one top-3 driver fired between 30 and 100 times; the top one or two are usually correct, but positions 3+ can shift on a re-run.

Insufficient means at least one top-3 driver fired fewer than 30 times; treat the rankings as indicative only, and re-run with more iterations (or untick “Stop early on convergence” on the Simulation page) before defending the result.

Why this matters: each driver’s ranking is essentially the mean of its impact across only the iterations where THAT risk fired AND THAT outcome landed in the tail. Below ~30 fires per top-3 driver, the mean is dominated by sampling noise. The chip surfaces this directly so you don’t have to do the arithmetic in your head. See §6.4 for the recommended early-stop policy and the rule-of-thumb for required iterations by driver probability.

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 responses on the events that contribute most to worst-case (and best-case) scenarios.

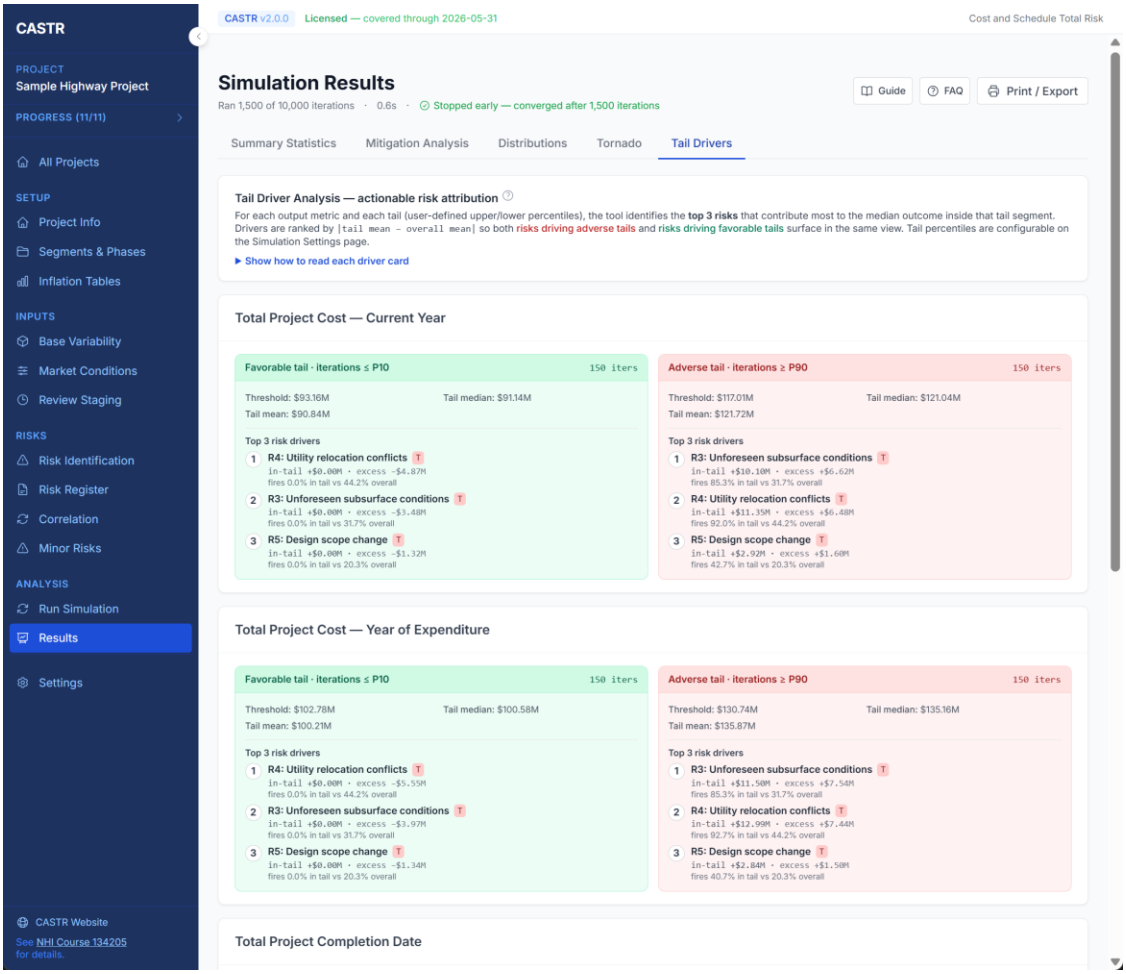

Figure 7.4 - Results Tail Driver Analysis

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

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 response 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.

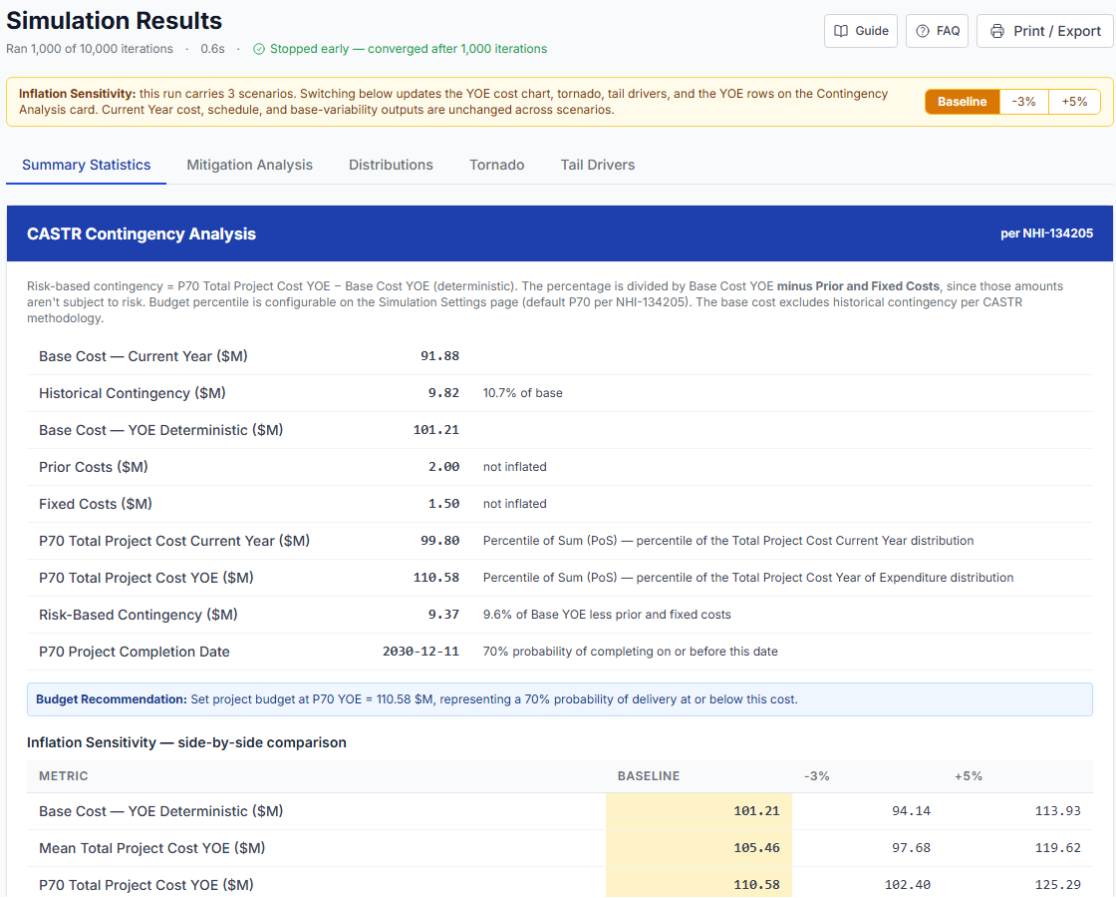

Figure 7.5 — Results Summary with the scenario chip selector (top), the CASTR Contingency Analysis card, and the side-by-side Inflation Sensitivity comparison table. Switching the chip selector repaints every other tab on this page accordingly.

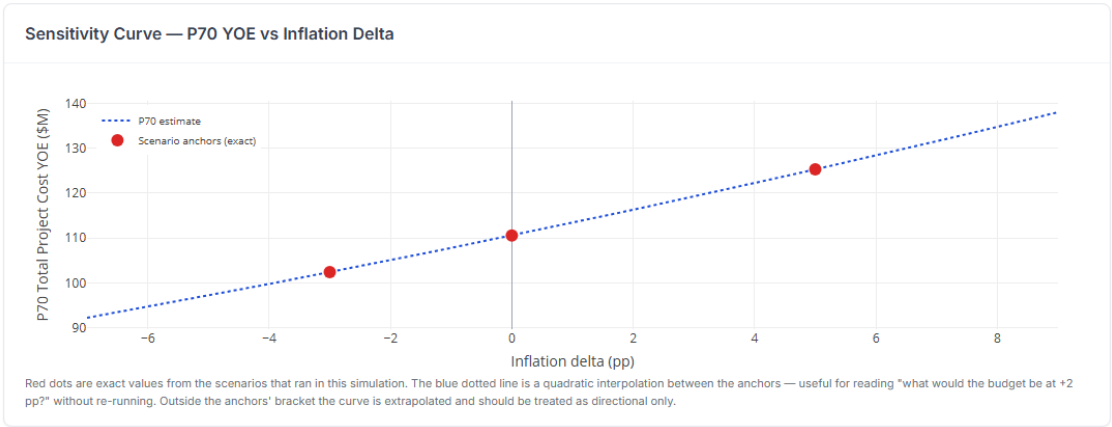

Figure 7.5a — Sensitivity Curve. Red dots are exact simulated values at each scenario delta; the dotted line is the quadratic interpolation between them. Outside the bracket the curve is extrapolated and should be treated as directional.

7.6 Word Export

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 Response 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.

7.7 Response Analysis

When the simulation Run Mode is set to Compare Both, the Results page exposes a Response Analysis tab with two side-by-side cards: the read-only Response Analysis card on the right (ROI summary table plus a 2x2 grid of pre-vs-post overlay charts for Cost CY, Cost YOE, Total Project Duration, and Total Project Completion Date) and a new Response What-If card on the left.

Response What-If panel. Lists every risk that has residual values defined (the only risks where RES can meaningfully change anything). Each row shows the risk's register number, name, residual probability, and which axes have residuals (cost / sched chips). Hover any row to see the full risk event name. The starting state has every candidate checked — the auto-baseline. Toggling a checkbox runs only the post-response Monte Carlo (debounced 350 ms; typically under a second) and updates the Response Analysis curves and ROI table live. The header reads Baseline, recomputing post-response pass..., or engine result (N/M on) so you always know which path produced the displayed numbers. Reset to baseline / All off helper buttons sit above the list. Toggles in this panel are LOCAL — they never write back to the Risk Register, so you can experiment freely. To commit a selection, mirror it on the Register RES column and re-run a saved simulation.

Resizable divider. A draggable vertical gutter between the two cards lets you rebalance the column widths to suit your screen. Drag to resize, double-click to reset to the default 320 px width. The chosen width is remembered across page reloads.

Total Project Duration — calendar time. Total Project Duration is reported as CALENDAR TIME — the per-iteration project completion date minus the earliest planned start, expressed in months. Duration and Completion Date move in lockstep — mitigating a critical-path risk shifts both metrics by the same amount, and a non-critical-path risk affects neither. Older simulation runs in the database may carry a legacy 'sum of all phase-

segment durations' value in their stored results JSON; re-run them to pick up the calendar-time definition.

Distribution color convention. All four distribution charts on the Distributions tab and in the Response Analysis pre-vs-post overlay (Cost CY, Cost YOE, Total Project Duration, Total Project Completion Date) share a unified color palette: light blue for histogram bins above the budget percentile, and your selected fill color (Simulation Settings → Distribution Display → Fill below budget) for bins at or below the budget cutoff.

Contingency envelope basis. The Implied % readings on the Contingency Envelope card (both on the Risk Identification Triage view and on the Risk Register) are computed against the Base Cost CY column from the Base Variability page (= Cost Estimate minus traditional contingency) instead of the Cost Estimate column itself. Using the pre-contingency base as the denominator avoids double-counting the sponsor's contingency in the implied % reading.

7.8 Joint Confidence Tab

The **Joint Confidence** tab — the third tab on the Results page, immediately after Distributions — plots the JOINT distribution of YOE Cost and Total Project Completion Date. It shows the two outcomes paired per-iteration rather than as independent marginals, and it is the chart you reach for when you have to defend a budget that satisfies a cost target AND a schedule target simultaneously.

Older saved runs without the paired cost/schedule samples show a "Joint plot unavailable" placeholder on this tab — re-run the simulation to generate the data the chart needs.

When the simulation does not produce an end-date distribution (rare — happens only when there are no base-uncertainty rows), the schedule axis automatically falls back to Total Project Duration in months.

What the chart shows

Three layers, drawn back-to-front:

1. **Scatter cloud.** Each blue dot is one Monte Carlo iteration plotted as (completion date, YOE cost). The cloud's overall shape tells you whether cost and schedule move together — a tight diagonal cloud means strong positive correlation (delays drive overruns); a round cloud means largely independent.
2. **Isoprobability contours.** Three contour lines at **50 % (red, densest core)**, **80 % (orange)**, and **90 % (yellow, outer)** mark the smallest joint regions that contain those fractions of iterations. Tightly-packed contours mean the distribution is sharply peaked; widely-spaced contours mean heavier tails. A miniature legend in the upper-left of the plot identifies the colors.
3. **Draggable P-budget crosshairs.** Two dashed black lines — one vertical (schedule), one horizontal (cost) — start at the marginal P-budget on each axis (the same P-

value used on the Contingency Analysis card). Each line carries a label that updates live as you drag it: “Schedule P{n}” rides along the top of the plot above the vertical line, and “Cost P{n}” rides along the right edge next to the horizontal line.

A green star marks the **deterministic Base point** at (Base Project Completion Date, Base Cost YOE). The base completion date is the maximum end-date across base-variant phase-segment rows from the Base Variability table; the base cost is the scenario-aware Base Cost YOE Deterministic. Hovering the star shows the exact pair. The base point is the no-uncertainty answer the Monte Carlo is sampling around — looking at how far up and to the right the dots spread from the star is the visual analog of the contingency conversation.

The four badges below the plot

- **Marginal P{budget} Cost** — the marginal P-value of cost (e.g. P70 = $903 M). Fixed reference; doesn’t change when you drag.
- **Marginal P{budget} Date** — the marginal P-value of completion date. Fixed.
- **Crosshair position** — the (cost, date) pair the lines are currently set to.
- **Joint coverage (live)** — the percentage of iterations that land below BOTH crosshairs, i.e. inside the rectangle to the lower-left of their intersection. This is the live readout that makes the chart interactive — drag a line and watch this number change.

Interpreting joint coverage

When you start at the marginal P70 × P70 crosshair position, joint coverage will almost always be LESS than 70 % (commonly 50–65 %). That is the central insight of the joint plot:

> Being below the marginal P70 on cost does not guarantee being below the marginal P70 on the date.

Cost and schedule are positively correlated on most projects (delays drive overruns), so a non-trivial fraction of iterations land below the marginal cost line but above the marginal date line, or vice-versa. To budget at a TRUE joint P70, you need to push BOTH lines outward. There is no single (cost, date) point that achieves joint P70 — there is an entire P70 contour curve, and you choose the trade-off:

- More cost headroom → tighter date allowance.
- More date headroom → tighter cost allowance.

The contour line at the budget percentile (which you can read off the orange line for joint P80, or interpolate near the red line for joint P50) traces every (cost, date) pair that delivers that joint coverage.

Why joint P70 × P70 is below 70 % — the algebra

Let `A` = "iteration is at or below P70 cost" and `B` = "iteration is at or below P70 date." By definition of the marginal percentile, `P(A) = 0.70` and `P(B) = 0.70`. The joint-coverage badge shows `P(A AND B)`. Three reference cases bracket the answer:

- **Perfectly positively correlated** → `P(A AND B) = 0.70` (every below-P70-cost iteration is also below-P70-date).
- **Statistically independent** → `P(A AND B) = 0.70 × 0.70 = 0.49`.
- **Perfectly negatively correlated** → `P(A AND B) = max(0, 0.70 + 0.70 − 1) = 0.40`.

A reading of, say, **0.607** sits comfortably above the 0.49 independence baseline and below the 0.70 perfect-correlation ceiling — which is the "positively but imperfectly correlated" regime that almost every infrastructure project lives in. Many of the same risks drive both overruns and delays, but they don't move in lockstep.

Why joint P70 × P70 is below 70 % — the visual

The two crosshairs slice the plot into four quadrants. Marginal P70 means **70 % of dots sit in the bottom row (below date)** AND **70 % sit in the left column (below cost)** — but those are two DIFFERENT 70%s.

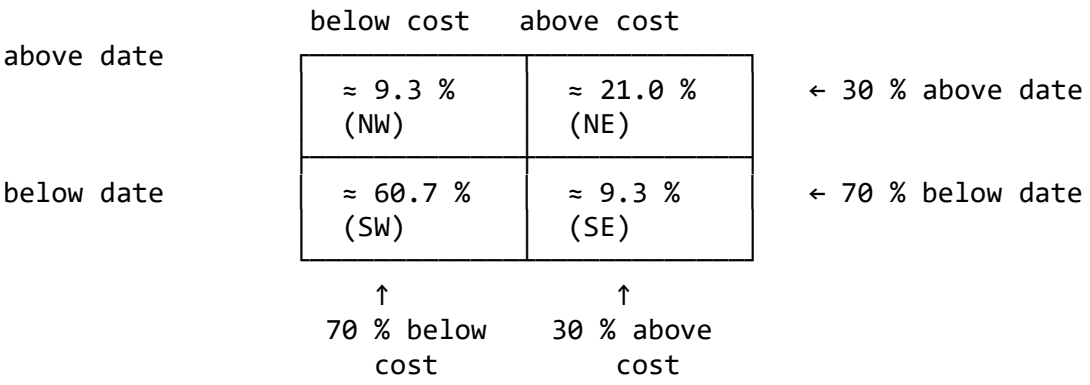

The SW quadrant is your joint coverage — both conditions met. The NW and SE quadrants capture iterations that "succeeded on one axis but failed on the other." With joint coverage at 60.7 % and marginals at 70 %, each of those off-diagonal quadrants holds about 9.3 % of iterations (= 70 % − 60.7 %). The NE quadrant — above both thresholds — holds the remaining 20.7 %.

Why this matters for budgeting

Citing the marginal P70 cost AND the marginal P70 date together as your "P70 budget" implicitly assumes those two events happen on the same iteration. They don't. There is only a ~60 % chance both targets are hit simultaneously. To honestly defend a TRUE joint P70 budget, drag both crosshairs outward until the live coverage badge reaches 70 %. The

contour line at P70 traces every (cost, date) pair that achieves it; you pick the trade-off you can defend.

Inclusion–exclusion sanity check: `P(A OR B) = 0.70 + 0.70 − 0.607 = 0.793` → 79.3 % of iterations land below at least one threshold; the remaining 20.7 % (the NE quadrant) are above BOTH — the "fail on both axes" region.

How to use the card

1. Read the Marginal P{budget} Cost and Marginal P{budget} Date badges first — these are your starting reference points.
2. Note the live Joint coverage badge at the marginal position. If it is well below the marginal P-value (e.g., joint coverage = 55 % at P70 × P70), you have meaningful positive cost-schedule correlation.
3. Drag the schedule line to the right until you reach a date you can defend; watch the joint coverage rise.
4. Drag the cost line up until you reach a budget you can defend; watch the joint coverage rise further.
5. Aim for a joint coverage that matches your agency's confidence target (often the same number as the marginal Budget Percentile).
6. Click **Reset to P{budget}** to snap both lines back to the marginal starting point.

Multi-scenario runs

When the run carries multiple inflation scenarios, the cost (Y) axis re-renders against the selected scenario's per-iteration cost samples. The schedule (X) axis is invariant. The scenario name appears in the card header.

Help icon

The header carries a question-mark icon that opens a popover summary plus a deep link to this manual section.

7.9 Inflation Scenario Comparison — YOE Cost Distributions

A side-by-side **violin + inner box** chart that sits at the bottom of the Distributions tab whenever the run carries two or more inflation scenarios. Each scenario gets one violin showing the YOE total-cost distribution shape; the inner box marks the P25 / median / P75 quartiles with whiskers extending to ±1.5×IQR.

Why this chart exists

The CDF overlays higher up on the Distributions tab can carry pre/post comparisons cleanly with two curves, but **three or more inflation scenarios overlapping on the same axes** become hard to read — the eye loses track of which curve belongs to which scenario, and shape differences disappear behind crossings. Violins separate the scenarios onto distinct slots of the x-axis so the reader can see at a glance:

- which scenario has the **highest median** (where the central case lands),
- which scenario is **widest** (most uncertainty),
- which scenario is **most skewed** (longer tail to the right means heavier upper risk under that inflation assumption).

How violin width is formed

Violin width at any y-value is proportional to a kernel density estimate (KDE) of the samples at that value — essentially a smoothed histogram, mirrored left and right of a center line:

- **Wide spots** = many simulated iterations cluster around that cost value (the mass of the distribution).
- **Narrow tapering** = few iterations there (the tails).
- **Symmetry** is purely visual; both sides of each violin show the same KDE.

Plotly normalizes each violin's maximum width to a constant by default, so a violin built from 100 samples and one built from 10,000 samples look the same width on the page. **Only the SHAPE of the violin communicates information; the absolute width does not.** If you need to compare absolute spread between scenarios, read the box quartiles, not the violin envelope.

Reading guide

Visual	Interpretation
Two violins of similar shape, one shifted vertically	Inflation moved the cost level but did not change the underlying risk shape — typical of broad-based inflation hitting all phase-segments roughly proportionally.
One violin much taller and narrower than the others	That scenario produced a tighter forecast (less uncertainty), often because the inflation assumption suppressed long-tail variance via lower compounding.
One violin with a noticeable bulge near the top	That scenario concentrates more mass in the upper cost band, indicating heavier upper-tail risk under that inflation rate.
All three violins look nearly identical	The project has very little inflation-sensitive cost (most spend in CY-anchored phase-segments) **or** the deltas you set are too narrow to move the engine meaningfully. Widen the deltas on the Run Simulation page to confirm the engine is responding.

When the chart hides

- Single-scenario runs (only the baseline ran) — no comparison value.

- Older runs that don't include per-scenario `cost_yoe_samples` arrays — re-run to populate them.
- Fewer than two scenarios contain sample data.

Where to go next

- **CDF view of the same data**: scroll up on the Distributions tab to the per-scenario chart strip — same scenarios, same axes, classical percentile-by-percentile reading.
- **Joint cost+date trade-off under each scenario**: the Joint Confidence tab also re-renders against the selected scenario chip.
- **Pre/post response-program comparison rather than scenario comparison**: see §10.8 in the Risk Response Analysis chapter.

Violin shape catalog — how to interpret what you see

Every violin's silhouette is a smoothed mirror of the YOE cost histogram for that scenario. Reading shape is at least as informative as reading percentile values, because two distributions with identical P25/P50/P75 can look very different. Six patterns cover most of what you'll encounter on a CASTR project:

1. Symmetric bell — fat in the middle, tapering equally to both tails. The Monte Carlo iterations cluster around a single typical cost with balanced upside/downside risk. Often appears when base variability dominates and the risk register is thin. Read the median (the line through the box) as a defensible central estimate; the box width tells you the typical-case spread.

2. Right-skewed teardrop — bulge low, long thin neck stretching upward. The canonical cost-distribution shape. Costs can blow up but only go down so far, so the upper tail is heavier than the lower. The MEDIAN sits noticeably below the MEAN; budgeting at the mean systematically under-funds the project. Use the box's upper whisker (P75 + 1.5×IQR) as a sanity check — anything beyond it is a candidate "tail-driver" risk worth investigating on the Tail Drivers tab.

3. Left-skewed teardrop — bulge high, long thin neck stretching downward (rare for cost). Almost always indicates an opportunity-heavy register or a downside-capped distribution (e.g., a fixed-price contract clause). The mean sits above the median. Cross-check the Risk Register for opportunities that may have unrealistic upper bounds.

4. Bimodal — two distinct bulges separated by a narrow waist. Two scenarios are competing inside one Monte Carlo. Almost always caused by a single high-impact, mid-probability risk: when it FIRES, iterations land in the upper bulge; when it DOESN'T, they land in the lower bulge. The narrow waist between them is the gap between "risk happened" and "risk avoided." Identify which risk on the Tornado tab dominates variance share — that's your bimodal driver. Bimodal cost distributions are honest but hard to defend; consider mitigating the dominant risk to collapse the upper bulge.

5. Tall and narrow — barely any horizontal taper, the violin reads almost like a vertical bar. The simulation produced a tight forecast. Two interpretations: (a) the project genuinely has low variance — small base-uncertainty percentages, well-understood risks, narrow distributions — or (b) **the risk register is under-coverage** and the engine has nothing to vary. Cross-check the Distribution Health card: a "Narrow" coverage chip + tall-narrow violins together usually means missing risks, not a confidently estimated project.

6. Asymmetric box, long whiskers — quartiles squashed to one side, whiskers reaching far on the other. The median is offset within the IQR. Almost always associated with strong skew (#2 or #3 above). The asymmetric box is a quick visual proxy for skewness: if the upper whisker is much longer than the lower, the distribution has a heavy upper tail.

Comparing scenarios on the same chart

When violins for multiple inflation scenarios share the y-axis:

- **Same shape, vertical shift** — inflation moved the cost level uniformly. The risk profile (variance, skew, tail behaviour) is unchanged across scenarios, only the absolute dollars changed.
- **Same baseline, taller / fatter at the top under higher inflation** — the high-inflation scenario amplifies upper-tail risks because YOE compounds with delay, and delayed iterations get hit twice (delay + inflation). Common pattern when schedule-impact risks are present.
- **Different shapes between scenarios** — usually means the scenario stresses a specific phase of the project unequally (e.g., one scenario backloads spend, where inflation has more compounding effect). Worth investigating on the per-phase-segment charts.

A scenario whose violin is markedly skinnier than the others under the SAME base risk register usually indicates the engine is producing less variance under that scenario — typically because lower inflation suppresses the compounding amplification of schedule risk. Higher inflation tends to widen the YOE cost distribution.

Common pitfalls when reading violins

- **Plotly normalizes max width per violin.** A "wide" violin and a "narrow" violin can both span the same horizontal envelope on screen even if one represents 10× more samples. Only use width to read SHAPE, not absolute spread. For absolute spread, use the box quartiles or the whiskers.
- **The whiskers are at ±1.5×IQR, NOT min/max.** Samples beyond the whiskers exist (visible as the violin's thin neck extending past the whisker line) but are conventionally treated as outliers. The min and max of the distribution can be far above the upper whisker.
- **Median ≠ mean.** Violins with strong right skew show the median (line through the box) noticeably below the bulge's "fat point" — that fat point is closer to the mode. The mean is even further to the right. Pick the right central tendency for your

audience: median for “what’s the typical project?”, mean for expected-value calculations.

7.10 Cost of Delay

The Cost of Delay card on the Summary tab answers a single workshop-defense question: *if every phase-segment’s midpoint slips by N years, how much more budget do we need?* Pick a delay (1–5 years) from the dropdown; the table fills in three column groups for each inflation scenario in the run.

Base Cost YOE (deterministic, exact). The deterministic Base Cost in YOE dollars at the planned dates and at the delayed dates. Computed by shifting both start and end dates of every phase-segment forward by N years and re-running the same per-segment inflation walk the simulator already uses. The reported delta is mathematically identical to what you would get by manually shifting every phase-segment’s start and end dates on the Base Variability page and re-running the simulation — no Monte Carlo needed for this column.

Risk-Based Contingency (approximate). Budget − Base on each side. Tells you how much of the budget growth is contingency creep on top of base inflation versus base inflation alone.

P{N} Budget Total YOE (approximate). The contingency-inclusive budget total at the project’s budget percentile from the live simulation, then approximated for the delayed case. Method: scale only the riskable portion of the planned budget percentile (everything except prior + fixed flat costs) by the same per-segment inflation ratio the deterministic Base Cost walk produced. Flat costs are held constant since inflation does not apply to them. The delta on this side is the dollar amount sponsors authorize against — *we need X more dollars in our budget if we slip N years*.

Decomposition. By construction, ΔBase + ΔContingency = ΔBudget on every row. Workshop leaders can point to the breakdown and explain *where* the additional budget ask is coming from — pure base-cost inflation, or contingency growth on top of it.

Why the Contingency and Budget columns are approximate, not exact.

- Schedule risks already perturb each phase-segment’s midpoint per iteration. The N-year uniform shift on top of those per-iteration shifts is not exactly the same as scaling the aggregate percentile by a single ratio. The error is usually small — the local inflation rate does not change much over a 1-month schedule shift — but for high-fidelity answers re-run the simulation with shifted dates.
- Non-uniform slips (e.g. PE delays but CN does not) are not modeled here. The card assumes every phase-segment slips by the same N years.
- The budget percentile read here uses the PoS method (top-level total cost percentile). If the project uses SoPSP or SoSP, the planned-side number may differ from the Contingency Analysis card by a few percent.

When to re-run instead. For a fully risk-aware, exact answer to a non-uniform slip — only some phase-segments delayed, or different segments delayed by different amounts — re-run the simulation with the relevant phase-segment start and end dates shifted on the Base Variability page. The Cost of Delay card answers the uniform-slip case quickly and inline; the re-run handles everything else.

Chapter 8: Data Import and Export

8.1 Risk Import from Excel

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.

Template columns. The Risks sheet of the import template carries the full set of risk attributes used by CASTR: Row #, Include (Y/N), Category ID, Segment / Phase, Risk Event Name, Risk Response Summary, Probability (0-1), Dependency Type, Dependent Risk Name, Cost Threat/Opp, Cost Dist Type, Cost P1–P3, Sched Threat/Opp, Sched Dist Type, Sched P1–P3, Sched Additive (Y/N), and the documentation block — Details: Description, Details: Response Actions, Details: Status / Assumptions, Details: Risk Trigger, Details: Risk Response Strategy, Details: Risk Owner, Details: Updated Date (YYYY-MM-DD), Details: Notes, plus the mitigation columns (Risk Response, Mitigation Cost, Mitigation Owner, Mitigation Due Date, Residual Probability, Residual Cost / Sched dist parameters). Every cell has a hover comment in the template explaining what it accepts; blank cells fall back to model defaults so you only need to fill in what is meaningful for the risk.

Starter Library import template. The Starter Library page has its own three-step import wizard (Import from Excel → preview → commit) for the app-global library. Its Risks sheet has eight columns: Cat ID, Risk Event Name, Risk Description, Primary Impact (Cost / Schedule / Both — dropdown), Typical Phase (comma-separated list of phases; “All” is a wildcard), and three complexity-tier flags (Non-Complex, Moderately Complex, Most Complex — X to mark applicability). The importer accepts “Default Categories” / “Default Risks” as sheet-name aliases of “Categories” / “Risks”, and detects an extra label column

between Cat ID and Risk Event Name by header name rather than position so workbooks authored elsewhere round-trip cleanly.

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 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 and large-batch operations automatically take their own checkpoints before running: Replace-mode import, Clear All Risks, Restore of a prior checkpoint, Bulk Delete, and Bulk Clone (both via 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 …", "Before bulk clone …", 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 RES 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.

8.5 Per-Risk-Owner Reports

Purpose.

After a workshop, every risk is typically assigned to a Risk Owner who is responsible for monitoring and executing the planned response. The full Risk Register is too wide and too complete for a per-owner hand-off. The By Owner feature on the Risk Register page slices the register by the Risk Owner field and produces a focused, professional report covering only the risks assigned to each owner — ready to email or hand off as a printable PDF or a Word document.

Opening the By Owner panel.

On the Risk Register page, click the By Owner button on the action toolbar (next to Print, CSV, and Excel). A popover opens listing every distinct Risk Owner found on the register, sorted alphabetically, with a count of risks assigned to each owner. Risks without an owner appear in a separate (Unassigned) row at the bottom — a reminder of the rows that still need an owner before the workshop closes out.

The Risk Owner field is the same one shown on each risk's Risk Details panel on the Risk Identification page (Owner field under Status & Assumptions). Edit the owner there, sync to the register, and the By Owner panel picks it up automatically.

Columns to include (saved per project).

The Columns to include checklist at the top of the popover controls exactly which fields appear in BOTH the Print handout and the Word document. Every register column is listed, grouped by section (Identification, Probability & Impact, Documentation, Status & Ownership, Response, Triage, Post-Simulation). Tick the columns you want each Risk Owner to receive and untick the rest. Use Select all to start with every documented field, or Clear to start from an empty list and pick exactly what each owner needs to act on.

The selection is saved to the project file the moment a checkbox is toggled — a small status indicator next to the count flips between 'saving…' and 'saved' so you can see the write happen. Reload the page, switch projects, or open the project in another session and the same columns are still ticked. Set it once per project and forget it. The on-screen Register column-picker (Brief / Standard / Full / Custom) is independent of the per-owner column list, so the analyst-facing register view and the owner-facing handout can have completely different column sets.

Print and Word actions.

Each owner row in the popover has a Print button and a Word button. Both render the same content (whatever columns are currently selected); they differ only in format.

Print opens a styled popup window optimised for the browser's Print / Save-as-PDF dialog. Each risk renders as a card with the owner's name in the document header and a key/value field table for every selected column. Use this when you want to email the owner a PDF, or print a paper hand-out directly.

Word downloads a .docx file with a title block, a summary table (owner name, risk count, generated date, source project), and one labelled sub-section per risk. The file is named ' - - Risk Assignment - YYYY-MM-DD.docx', drops cleanly into the project documentation folder, and can be edited or annotated before being sent on.

All owners — one Word doc each.

When the project has at least two assigned owners, a footer button labelled ‘All owners — one Word doc each (N files)’ appears at the bottom of the popover. One click downloads a separate .docx per owner, staggered slightly so the browser does not drop any of the downloads. Each file follows the same naming convention as the per-owner Word button, so the files sort cleanly in the Downloads folder for easy distribution.

Notes.

If every column is unchecked, both Print and Word are dimmed — there is nothing for the report to show. Tick at least one column (or click Select all) to re-enable them. If the browser blocks the Print popup, click the pop-up icon in the address bar and allow popups for the CASTR site, then click Print again. CASTR currently supports a single Risk Owner per risk; if two people share responsibility, list the primary owner there and capture the secondary owner in the Response Owner field or the project’s assumptions text.

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 curated set of 25 highway-construction risk categories (Scope, Schedule, Estimate, Design, Geotechnical, Site Conditions, Bridge & Structures, Pavement, Hydraulic / Drainage / Stormwater, Environmental, Cultural & Historic, Right-of-Way, Utility, Railroad, Third-Party, Public Involvement, Procurement & Contract, Construction, Quality / QA-QC, Safety / Health, ITS / Technology, Market / Escalation, Funding, Organizational, External / Force Majeure) and 168 risk events synthesized from FHWA Major Projects guidance, Caltrans PRMHB, WSDOT CEVP/CRA, FDOT PMG 405, GDOT PDP, and NCHRP 08-60.

Each starter risk carries:

- **Cat ID** — the category code (e.g. `ENV`, `ROW`, `STR`).
- **Risk Event Name** — short label shown in the picker and copied onto the project risk row.
- **Risk Description** — longer narrative copied into the risk's Details on insert.
- **Primary Impact** — `Cost`, `Schedule`, or `Both`. Indicates which axis the risk most affects.
- **Typical Phase** — one or more of Planning, PE/Design, ROW, Environmental, Procurement, Construction. A risk may apply to several phases; `All` is a wildcard for every phase.
- **Project Complexity Tiers** — three independent flags (Non-Complex / Moderately Complex / Most Complex) marking which project tiers the risk typically applies to. A risk may apply to one, two, or all three tiers.

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 — Cat ID (1–4 uppercase letters), Description, and a read-only count of how many risks reference the category. Sortable by Cat ID or Description (click the column header to toggle asc / desc / no-sort).

Risks — every column is inline-editable:

- **Cat ID** — dropdown of the current categories.
- **Risk Event Name** and **Description** — free text.
- **Impact** — single-select dropdown (Cost / Schedule / Both).
- **Phase** — multi-select popover; tick any combination of the six phase tokens, or tick "All" (mutually exclusive with the specific phases) to indicate the risk applies in every phase.
- **NC / M / X** — three checkboxes for the complexity tiers (Non-Complex / Moderately Complex / Most Complex).

A chip bar above the table filters risks by one or more categories (counts update live). All fields save on blur or, for checkboxes and the phase popover, on toggle. Blank required fields are back-shaded light yellow as a reminder.

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. 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.

Every column header carries a small help icon (?). Click it for a popover explanation of what the column drives downstream in the picker and Excel template.

Excel Import. 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 pre-filled with the current library so you can edit in place. The Risks sheet has eight columns: Cat ID, Risk Event Name, Risk Description, Primary Impact, Typical Phase, Non-Complex, Moderately Complex, Most Complex. The Primary Impact column has a dropdown enforcing `Cost / Schedule / Both`; the Typical Phase column accepts a comma-separated list (e.g. `Planning, PE/Design`) so a single risk can apply to multiple phases. The tier columns accept `X` (or `Y` / `true` / `1`) to mark applicability. 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.

The importer also accepts sheet names "Default Categories" / "Default Risks" as aliases of "Categories" / "Risks", so workbooks authored elsewhere round-trip without renaming sheets. Unknown sheets (e.g. "Read Me", "Project Complexity") are ignored. If the Risks sheet has an extra label column between Cat ID and Risk Event Name (some external references use one), the importer detects it by header name rather than position.

Restore Defaults. Each section has its own restore button: "Restore Category Defaults" beside the Categories header and "Restore Risk Defaults" beside the Risks header. Each wipes only its table and re-seeds it from the factory defaults that ship with CASTR; the

other table is not touched. This means you can refresh categories without losing custom risks, or refresh risks without losing custom categories.

When the contents of either table differ from the factory defaults, that section's Restore button turns amber and a * is appended to its label — a visual cue that you have unsaved deviations and that clicking Restore would discard them. The button reverts to the muted style once the table matches defaults again. Both buttons prompt for confirmation before wiping, and both note the "orphan" implications (restoring just risks may leave references to categories the user has since renamed/deleted; restoring just categories may strand risks that reference now-removed cat_ids — the picker still shows orphan risks but they appear under no category filter).

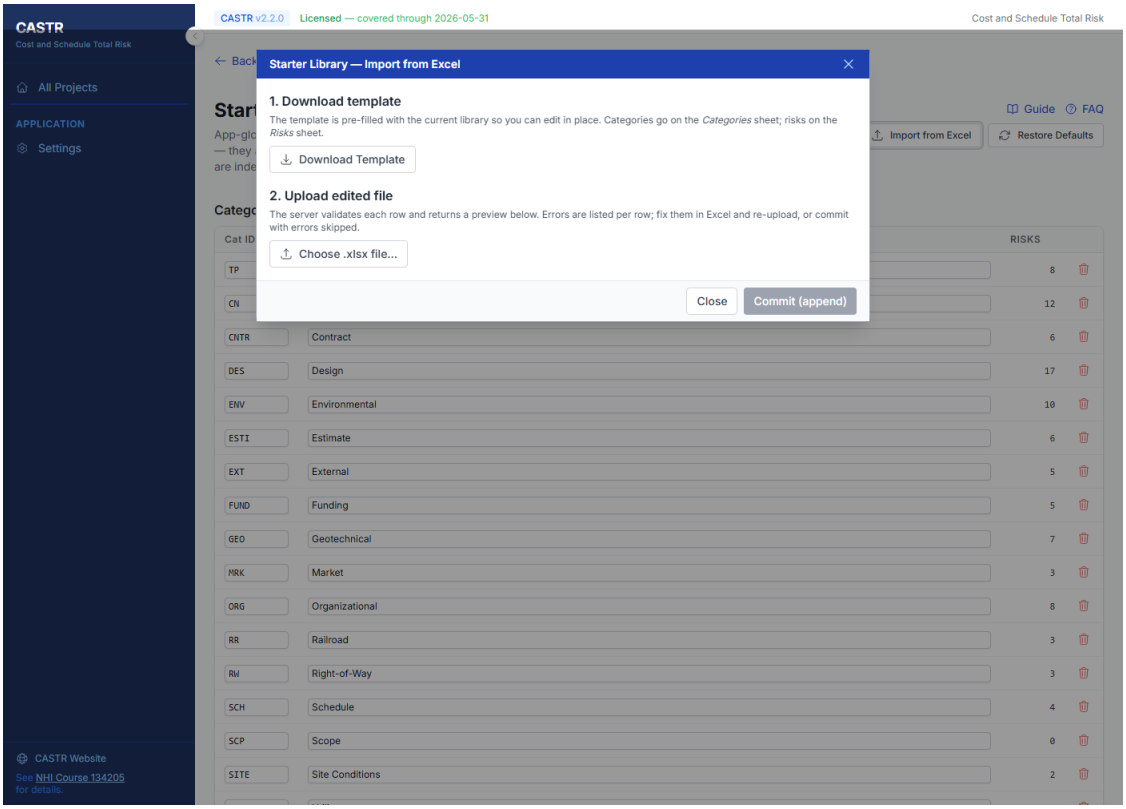

Figure 9.1 — Starter Library Import-from-Excel wizard. Step 1 downloads a pre-filled template; step 2 uploads the edited workbook and shows a preview; step 3 commits the changes.

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 the full library — every starter risk visible by default. You narrow the list with four independent filter mechanisms that compose with AND logic:

1. **Category rail (left column)** — every starter category with a live count of risks in it. Click a category to toggle it; select multiple categories with OR logic among them, or click “All categories” to clear the selection. A green checkmark on a category row means that category already exists on the project.
2. **Project Complexity** — three pill checkboxes (Non-Complex / Moderately Complex / Most Complex). Pick the tier (or tiers) that match your project; only risks marked applicable to a checked tier remain. Empty selection means no complexity filter.
3. **Primary Impact** — three pills (Cost / Schedule / Both). Use this to focus the picker on, say, cost-only risks for an estimate review.
4. **Typical Phase** — one pill per phase present in the library (Planning, PE/Design, ROW, Environmental, Procurement, Construction, All). A risk passes the phase filter if any of its phases matches a checked pill; risks tagged `All` always pass any phase filter (they’re true wildcards).

A search box at the top of the picker matches risk name, description, Cat ID, Primary Impact, and Typical Phase strings — handy for quickly finding “asphalt” or “permit” across the whole catalog. A “Clear filters” link appears once any filter is active.

Each row in the right-hand list shows the risk’s Cat ID, name, description, and a metadata strip with three colored chips: Impact, Phase (joined with · when multiple), and Tiers (the letters `NC`, `M`, `X` of the tiers that apply).

Tick the checkbox on any number of rows, 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 toast 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. The Impact / Phase / Tier metadata stays on the starter library row and is *not* copied onto the project risk — it’s a property of the library entry, used only for picker filtering, while the project risk has its own per-phase assignment that you set on the Standard view.

Chapter 10: Risk Response Analysis

10.1 Overview

CASTR provides a full response workflow that lets you capture the planned response to each risk, the cost of the response action, the residual probability and impact after the action is applied, and then run the Monte Carlo simulation in either pre-response, post-response, or compare-both mode to quantify the return on the response program. The workshop-style Residual Distribution Editor opens with a green backshade and an amber "RES" 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 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 bundled demo project ships with multiple fully-mitigated risks (with RES toggled on), so a new installation can be simulated end-to-end in Pre, Post, and Compare-Both modes out of the box.

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

10.2 Defining a Response on a Risk

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

On the Risk Identification table, click the View (or Res) 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 Response 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 Response Details section appears with the following fields:

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

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

Response Owner — the person or role responsible for executing the response action.

Response Due Date — when the response action must be complete. The Risk Register highlights past-due dates when RES 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-response 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 “Res”, 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-response 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 Response Details section — while the Response 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-response) PDF and CDF in red, and the residual (post-response) 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 Response Status

The Risk Identification filter bar includes a three-way segmented control: All / Res / Non-Res. 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 response plan.

10.4 Sync to Risk Register and the RES 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 response fields are carried forward via the sync.

The Risk Register gains an RES checkbox column positioned after Sched Dist. and before Status. This column, along with the other response columns (Response, P(RESID), Response Cost, Response Owner, Response 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.

RES is the master switch. It tells the Monte Carlo simulator, for this specific risk, to use the residual (post-response) parameters instead of the original ones when the run mode is Post-Response or Compare Both. A risk can be planned for a response (Risk Response = Mitigate, residual params filled in) without RES being checked — the RES checkbox is the separate commitment step that says “this response 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 RES, re-syncing never clears it. On the first sync of a row, RES is pre-checked if Risk Response = Mitigate, so the common case needs no extra clicks.

Past-due highlight — when RES is active and the Response 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 RES 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-Response (Inherent Risk) — the default. Uses P(occur) and the original cost/schedule distributions for every risk. RES is ignored; results are byte-for-byte identical to historical CASTR behavior.

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

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

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

10.6 Response ROI Summary and Charts

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

The tab contains two sections: a Response 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.

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

Reg # — register row number.

Risk Event — the risk name.

Response Cost — the certain cost of the response action.

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

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

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

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

A totals row at the bottom aggregates Response Cost, Pre-RES EV, Post-RES EV, and Net Benefit, with an overall ROI = Total Net Benefit / Total Response 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-Response, green is Post-Response, 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 response 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.

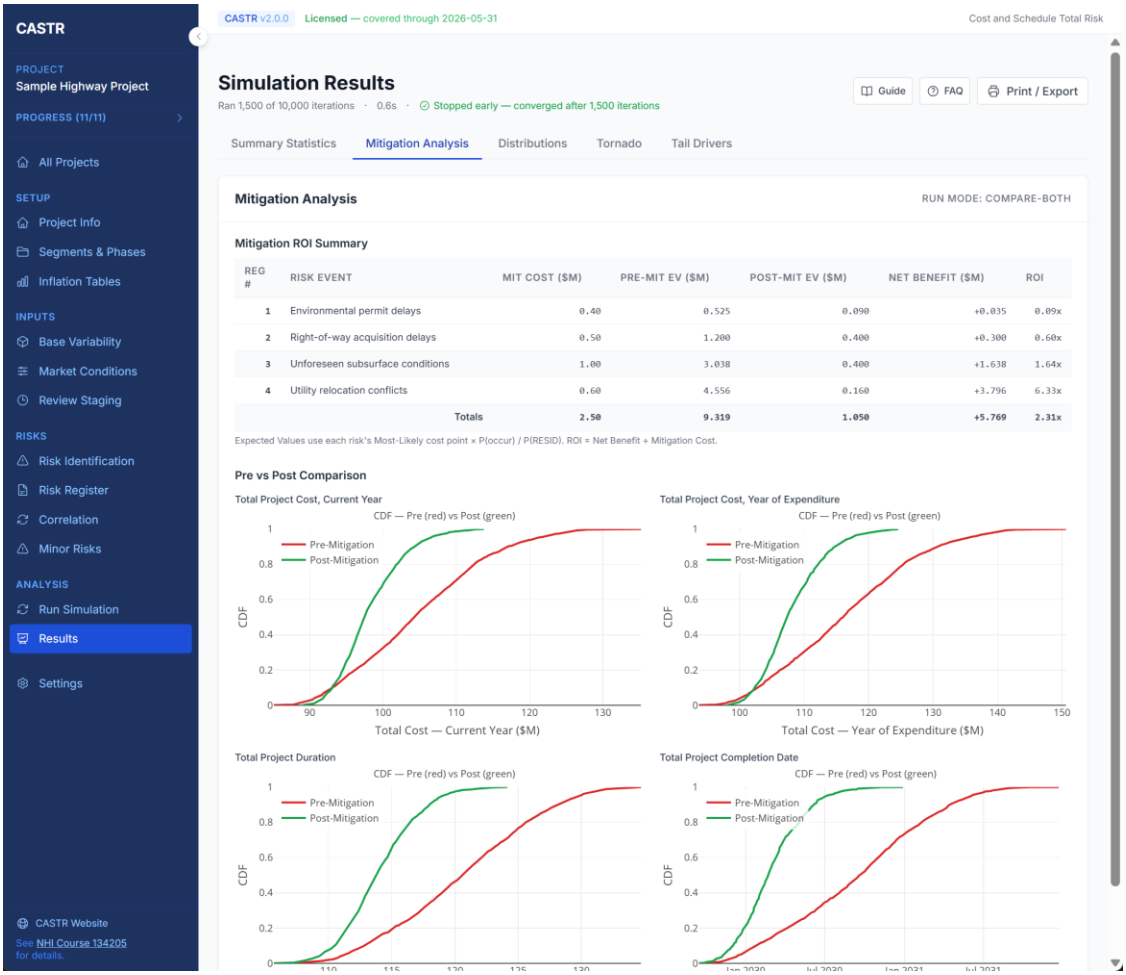

Figure 10.1 - Response Analysis: pre- vs post-response comparison of the project cost distribution, with response cost and ROI per risk summarized below.

Response ROI — by Inflation Scenario. When the run includes Inflation Sensitivity Scenarios (§6.1) AND the Risk Register has at least one active response, a second table titled “Response ROI — by Inflation Scenario (YOE)” appears below the headline CY ROI table. It reports Pre-RES EV, Post-RES EV, Net Reduction, and ROI per risk under each scenario’s per-risk inflation factor in YOE dollars. Response Cost is NOT scaled by inflation — response 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 response:

Step 1 — On Risk Identification, open a risk’s Details panel and set Risk Response to Mitigate. Fill in Response Cost, Owner, Due Date, P(RESID), and click the two Edit Residual

Distribution buttons to tune the post-response curves. The Details button for this row turns green.

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

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

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

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

10.8 Pre/Post Response Cost Distribution Shape

A **violin + inner box** chart on the Response Analysis tab that compares Pre-Response (red) and Post-Response (green) YOE total-cost distributions side-by-side. Each violin shows the smoothed shape of the cost distribution; the inner box overlays P25 / median / P75 with whiskers extending to ±1.5×IQR.

Why this chart exists

The CDF overlays in the Pre vs Post Comparison block immediately above the violins answer **"did the response program move the curve?"** — which is what budget percentiles need. But CDFs **flatten the shape** of the distribution: two distributions with identical P25 / P50 / P75 percentiles can have very different underlying shapes (bimodal vs unimodal, gentle vs steep skew), and the CDF compresses that into a single curve.

Violins keep the shape visible. Use both:

- **CDFs** for percentile readouts and the "what is the P70 cost?" question.
- **Violins** for distribution-shape diagnosis — "did the program reduce variance, or just shift the mean?".

Reading guide

The most leveraged response programs produce a post (green) violin that is **both** lower (median pulled down) **and** narrower (uncertainty compressed). Other patterns indicate different things:

Visual	Interpretation
Green violin shifted down, same shape and width as red	Response program reduced expected value but not variance. Common with high-Response-Cost / low-tail-impact responses — every iteration loses the same amount of cost regardless of which risks fired.

Visual	Interpretation
Green violin narrower at the top but same width at the bottom	Response program suppressed upper-tail risk specifically — the extreme overruns are gone, but the typical case is unchanged. Often the right pattern after a focused mitigation campaign on the highest-impact risks.
Green violin same shape as red but ONLY a small downward shift	The committed Response Costs are large relative to the EV reduction. Confirm against the Response ROI table — Net Benefit will be near zero.
Green violin lower AND narrower AND less skewed	The premier outcome: responses pulled the median down, compressed uncertainty, and removed the long right tail. Rare without targeted mitigations on multiple top tornado drivers.
Green violin nearly identical to red	Either too few risks have RES ON, residual distributions weren't actually tightened, or responses target non-critical-path schedule risks without predecessor links wired. See §10.5 and the workshop FAQ for diagnosis.

How violin width is formed

Width at any y-value is proportional to a kernel density estimate of the samples there. Wide = many iterations, narrow = few. Plotly normalizes max width per violin so a heavy-mass scenario and a light-mass scenario both span the same horizontal envelope — **only shape carries information**, not absolute width.

When the chart hides

- Pre-only or post-only runs — the chart requires both sides to render.
- Runs where the Joint Confidence card data wasn't shipped — same `joint_data.cost_yoe_samples` source on each side.
- Compare-both runs where one side returned zero samples (rare; would also break Joint Confidence).

Color convention

The pre/post color pair is consistent across the entire Response Analysis tab and the rest of the application: **red for Pre-Response, green for Post-Response**. Same convention as the CDF overlays directly above the violins.

Where to go next

- **Inflation-scenario comparison instead of pre/post**: see §7.9 in the Understanding Results chapter — same kind of violin chart, different comparison axis.
- **Per-risk ROI**: scroll up to the Response Analysis card's ROI table for the dollar-by-dollar accounting of which response saved how much.
- **Live what-if exploration**: the Response What-If panel next to the Response Analysis card lets you toggle individual responses on and off without re-running the full simulation; the violin will not auto-update on what-if toggles (it ties to the saved baseline) — re-run a saved simulation to refresh the violin.

Violin shape catalog — pre vs post pattern recognition

The interpretation table earlier in this section covers the most common pre-vs-post pairings. Below is a more thorough catalogue of shapes you may see on each side, plus what it means when the post (green) violin transforms into a different shape than the pre (red).

Single-violin shape interpretation

These apply to either side individually — pre OR post — and are the same shapes you'd read on the inflation-scenario violin chart (§7.9):

- **Symmetric bell** — balanced upside/downside. Often pre-response when base variability dominates.
- **Right-skewed teardrop** — canonical cost shape; long upper tail. Median below mean.
- **Left-skewed teardrop** — opportunities outweigh threats, or a hard upper cap exists. Rare for cost.
- **Bimodal** — one big risk dominates and either fires or doesn't. Each bulge is one outcome state.
- **Tall and narrow** — tight forecast (or under-covered risk register, see Distribution Health card).
- **Asymmetric box with long whiskers** — strong skew; the visible asymmetry is a proxy for the skewness number on the Summary tab.

Shape transitions — pre → post

Pre (red)	Post (green)	Interpretation
Right-skewed teardrop	Symmetric bell	Responses removed the upper tail. The dominant tail-driving risks were either avoided or had their residual upper bounds significantly tightened. **Best-case outcome** for a focused

Pre (red)	Post (green)	Interpretation
		mitigation campaign.
Bimodal	Unimodal (single bulge)	The dominant risk that was creating the upper bulge has been mitigated, transferred, or avoided. Iterations no longer split into “fired vs didn’t fire” outcomes. **Strong evidence the response program targeted the right risk.**
Bimodal	Bimodal, lower bulge larger	The dominant risk fires less often (probability dropped) but still has its original impact when it does fire. The lower bulge grew (more iterations where risk didn’t fire), the upper bulge stayed where it was. Verify the residual probability matches the workshop’s intent.
Bimodal	Bimodal, upper bulge moved down	The dominant risk’s IMPACT was reduced but not its probability. Both bulges still appear in roughly equal proportions; the upper one is just at a lower cost level. Workshop chose to “cap” the risk rather than “prevent” it.
Right-skewed	Right-skewed, shifted down	Responses bought down expected value uniformly across iterations — typical of broad insurance-style mitigations or contingency reserves. Median moved but variance is unchanged. **Net Benefit may be small relative to Response Cost.**
Right-skewed	Right-skewed, narrower	Responses compressed the tail without changing the median. Upper tail risks

Pre (red)	Post (green)	Interpretation
		were specifically targeted; the typical-case stays unchanged but the worst-case is bounded. **Often the best ROI per dollar spent.**
Tall and narrow	Tall and narrow, slightly lower	Either there was little room to improve (project was already tight) OR the responses are immaterial — the residual distributions are barely different from the originals. Cross-check the Risk Register that residual values were actually tightened.
Right-skewed	Same shape, identical position	**Red flag**: nothing changed. Common causes — too few risks have RES ON, residual values weren't tightened, or all responses target schedule-only risks while the chart shows cost. See §10.5 for the diagnostic checklist.
Right-skewed, wide	Left-skewed	Highly unusual. Indicates an opportunity-heavy response program that overshot — the post-response register is now downside-skewed. Can happen when residual distributions for opportunities were set too aggressively. Sanity-check the residual params for opportunities.

Diagnostic flow

When the pre/post comparison doesn't look like what the workshop expected:

1. **Confirm both violins are present.** Pre-only or post-only renders mean the run mode wasn't compare_both. Re-run from the Simulation page with Run Mode = Compare Both.

2. **Cross-check the median lines.** If the green median is at or above the red median, the response program is not reducing expected cost — almost certainly a configuration issue (residuals not tightened, or RES toggles off).
3. **Look at the box width, not just the whiskers.** Box width = IQR = the middle 50% of iterations. If the green box is the same width as the red box, the response did not reduce typical-case uncertainty even if it shifted the median.
4. **Check the long thin tails.** A residual long upper tail on the green violin after responses were applied means at least one major risk was NOT addressed — open the Tail Drivers tab and identify it.
5. **Verify against the ROI table.** The Response ROI Summary table immediately above the violins gives per-risk Net Benefit. If Net Benefit is dominated by one or two risks, the violin shape change should be driven primarily by those risks. If it isn't, something is mis-configured.

Cross-reference

The same violin shape catalog applies to the inflation-scenario chart on the Distributions tab — see §7.9 for the per-scenario interpretation guide. The mechanics of how violin width is formed (kernel density estimate, max-width normalization, whisker conventions) are covered there in detail and aren't repeated here.

10.9 Budget-Constrained Response Bundle

A budget slider that picks the highest Return-on-Investment (ROI) subset of your defined responses that fits a chosen total spend. Sits below the Response ROI Summary on the Response Analysis tab, collapsed by default — click to expand.

Why this card exists

The Response ROI Summary tells you each individual response's economics. The Budget-Constrained Response Bundle answers the next, harder question: **"If I only have $X to spend on responses, which combination should I fund?"** This is the classic knapsack problem (each response has a Cost and a Net Benefit; total cost must not exceed the budget; maximize total Net Benefit).

How the solver works

1. **Auto-include free wins.** Any response with zero Response Cost and positive Net Benefit is included unconditionally — those are free improvements with no budget impact.
2. **Skip negative responses.** Any response whose Net Benefit is less than or equal to zero is excluded from consideration regardless of budget — they are losing investments.
3. **Greedy ratio fill.** Remaining responses are sorted by Net Benefit divided by Response Cost (Marginal ROI) in descending order. The solver fills greedily: if the next response's cost fits the remaining budget, include it; otherwise skip it.

4. **One-step swap refinement.** After greedy fill, the solver checks whether replacing one already-selected response with one skipped response of higher Net Benefit would improve the bundle without breaking the budget. If yes, it swaps. This catches the common knapsack edge case where the last greedy item blocked a higher-value candidate.

The result is provably within (1 - 1/e) of the optimal bundle for the general knapsack, and effectively optimal when individual response costs are small relative to the budget — which is the typical CASTR project shape.

How to use the card

1. **Slider defaults to your total mitigation cost.** All affordable, positive-Net-Benefit responses are pre-selected. The bundle equals "everything you have already defined."
2. **Drag the slider down** to a constrained budget. Watch the Bundle Totals row at the bottom — it shows the cost of the selected subset, the total Net Benefit, and the resulting ROI.
3. **Read the Cumulative Net Benefit vs Cost curve** at the bottom. Each step adds the next-best response by ROI ratio. **The curve flattens where additional spend stops earning meaningful returns** — that is the natural cap on how much your response program is worth funding. The green-shaded region marks what the current budget actually buys.
4. **Compare the unspent dollar number** below the table. If you cap the budget at $5M but the optimal bundle costs only $4.7M, the remaining $0.3M is too small to fit any unfunded response and is genuinely unspent.

Reading the bundle table

Column	Meaning
Reg #	Risk Register number for the response — same key used everywhere on the application.
Risk Event	The risk the response addresses.
Response Cost (in millions of dollars, $M)\|Certaindollarscommittedtofundthisres$ M)	Pre-Response Expected Value (EV) minus Post-Response EV minus Response Cost — what funding this response actually nets you in current-year (CY) dollars.
Marginal ROI	Net Benefit divided by Response Cost. The greedy sort key. Free responses (Cost = 0) appear with infinity at the top.

The Bundle Totals row sums Cost and Net Benefit across the selected subset, then divides for the bundle-wide ROI. The footer line below the table reports candidate count and how many responses were excluded for non-positive Net Benefit.

When this card hides

When no risks on the register have responses defined — there is nothing to optimize and the entire Response Analysis tab is hidden too.

Important assumptions

- **Independence assumption.** The solver treats each response as independent of the others. If two responses address the same risk and **only one can be funded** (e.g., "negotiate ROW differently" vs. "condemn the parcels"), exclude the dominated option from the Risk Register before reading the bundle. Otherwise the solver may pick both and overstate Net Benefit.
- **Per-risk responses, not portfolio responses.** Each row is one risk plus one response option. The bundle does NOT model substitute responses on the same risk; that is a Risk Register modeling choice, not an optimization concern.
- **Static cost reduction estimate.** Net Benefit comes from the simulation's pre/post Expected Value calculation. If the workshop revisits the residual distributions for any risk, re-run the simulation; the bundle then re-optimizes against the new economics on the next page load.

Where to go next

- **Per-risk economics** — see the Response ROI Summary table immediately above the bundle (section 10.6).
- **Distribution-shape impact of the bundle** — switch off responses you have chosen to defer in the Response What-If panel and watch the Pre/Post curves recompute live (section 10.7).
- **Detailed reading of bundle vs. Pareto frontier** — the cumulative-benefit curve below the bundle table is your visual Pareto curve; the slope at any point is the marginal ROI of the next dollar.

10.10 Pre vs Post Comparison Charts

A 2x2 grid of overlay charts that sits below the Response ROI Summary and Budget-Constrained Response Bundle on the Response Analysis tab, comparing every iteration of the project **before responses (red, Pre-Response)** to **after (green, Post-Response)**. Visible only on compare-both runs (Run Mode = Compare Both on the Run Simulation page).

What the four charts show

Chart	What it answers
Total Project Cost — Current Year ($M) | "Did the response program reduce cost in today's dollars?" — the unleavened cost view, no inflation. | | Total Project Cost —	"Did the response program reduce cost in future-year (YOE) dollars?" — the budget-defense view. Diverges from the current-year (CY) chart whenever responses also

Chart	What it answers
Year of Expenditure ($M)	shorten schedule, since shorter schedule = less inflation compounding.
Total Project Duration (months)	“Did the response program shorten the project?” — calendar months from earliest planned start to the project completion date. Reads the critical path through the predecessor graph.
Total Project Completion Date	“When does the project actually finish?” — the absolute date axis version of Total Project Duration.

Reading each chart

Every chart is a Cumulative Distribution Function (CDF — probability of being at or below any given value). The y-axis runs 0 to 1 (probability); the x-axis is cost or schedule. Two curves are drawn:

- **Red — Pre-Response** — the Monte Carlo distribution computed with every risk’s original probability and impact distributions.
- **Green — Post-Response** — the same Monte Carlo, replaying the same random seeds, but with each response-active risk swapped to its residual probability and residual impact distributions, plus the Response Cost added as a certain commitment to the project base.

If the response program is reducing risk, **the green curve sits to the LEFT of the red curve** (lower cost, earlier date). If the curves overlap, the response program is not moving the needle — see section 10.5 for the diagnostic checklist.

Vertical dashed lines and the delta caption

When the Show Budget Percentile flag is on (Simulation page → Distribution Display → Show Budget Percentile), each chart adds two vertical dashed lines — red at the Pre-Response value at that percentile, green at the Post-Response value — and a delta caption below the chart reporting the difference at the budget percentile. This is the **headline number for the response program’s impact** at your agency’s confidence level.

For example:

- Pre P70: $868.79M
- Post P70: $740.58M
- Delta at P70: minus $128.22M

The minus sign means the responses moved the budget DOWN by $128.22M at the P70 confidence level. This is what you defend in front of executives.

Color convention

The pre/post color pair is consistent across the entire application: **red for Pre-Response, green for Post-Response**. The same palette is used on the Pre/Post Response Cost Distribution Shape violin chart (section 10.8) immediately below this 2x2 grid.

Sum-of-Percentiles caveat (PoS, Sum of Segment Percentiles, Sum of Phase-Segment Percentiles)

The cost distributions shown in this 2x2 grid are always the **Percentile of Sum (PoS)** total — the per-iteration sum of all phase-segment costs. If the run is configured for **Sum of Segment Percentiles (SoSP)** or **Sum of Phase-Segment Percentiles (SoPSP)** budget percentile method, the budget contingency on the Summary tab uses that reconstructed total, but these overlay charts still reflect the underlying PoS Monte Carlo distribution. A blue informational note at the top of the comparison block flags this on every compare-both run.

When the charts hide

- Pre-only or post-only runs — the grid only renders on compare-both runs.
- Single-axis runs (e.g., Cost only, no Schedule) — the schedule chart slot shows the available data; the missing chart leaves a gap.

Where to go next

- **Distribution shape diagnosis** — the violin chart immediately below this grid (section 10.8) shows the same data with shape information that CDFs flatten.
- **Per-risk economics** — the Response ROI Summary above this grid (section 10.6) shows which risks contributed which fraction of the headline delta.
- **Live what-if exploration** — the Response What-If panel (section 10.7) toggles individual responses on and off and recomputes the Post-Response curves under one second; useful when defending a partial response program.

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. 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: a 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: 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

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.

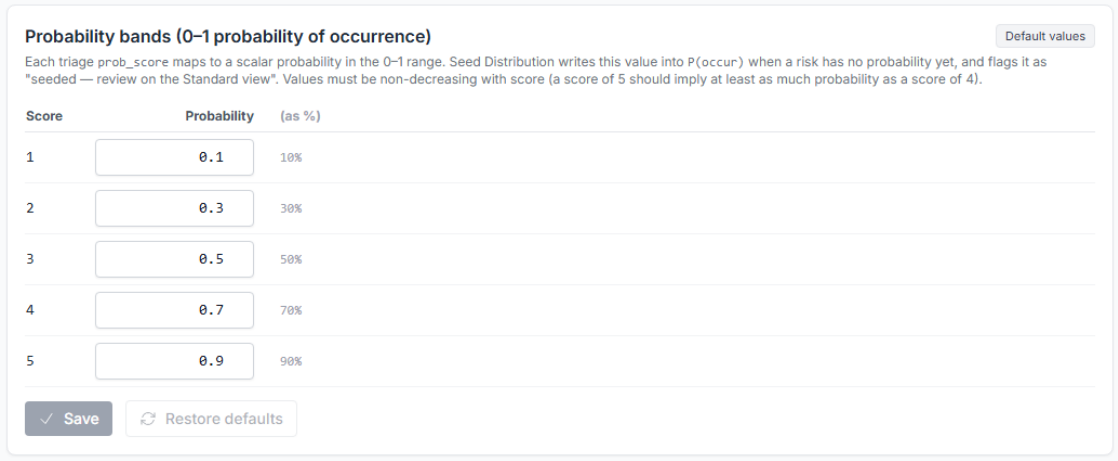

Figure 11.1 — Probability Bands editor. Each Triage probability score (1–5) maps to a scalar probability in the 0–1 range. Edits are global; not retroactive.

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.

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

A coordinated set of analytical tools helps 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

On the Risk Identification page in Triage view, the Contingency Envelope card has a three-way risk count column (Active / Reviewed / Inactive), two low-side State flags, and sortable headers across every column. The Triage envelope 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 the 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.

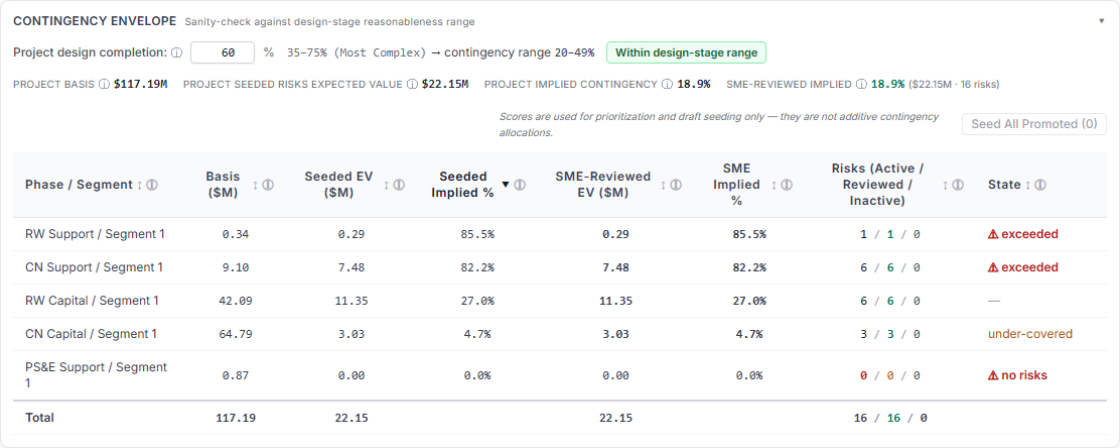

Phase / Segment	Basis ($M)	Seeded EV ($M)	Seeded Implied %	SME-Reviewed EV ($M)	SME Implied %	Risks (Active / Reviewed / Inactive)	State
RW Support / Segment 1	0.34	0.29	85.5%	0.29	85.5%	1 / 1 / 0	⚠ exceeded
CN Support / Segment 1	9.10	7.48	82.2%	7.48	82.2%	6 / 6 / 0	⚠ exceeded
RW Capital / Segment 1	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. Six sortable columns — Phase / Segment, Basis, Seeded EV, Seeded Implied %, Risks (Active / Reviewed / Inactive), and State. The SME-Reviewed columns are on the Standard-view and Risk Register envelope cards. 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

The State column flags both high-side and low-side conditions. High-side: approaching / exceeded (the implied contingency is approaching or above the design-stage band). Low-side:

• ⚠ 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 is sortable too, 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

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.

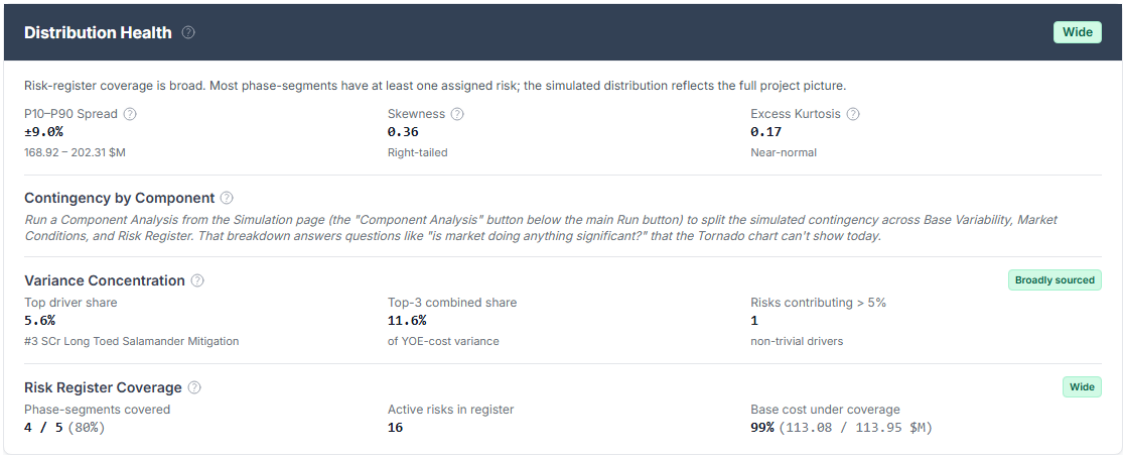

Figure 12.2: Distribution Health card showing the overall traffic-light, the three diagnostic cells (P10–P90 Spread, Skewness, Excess Kurtosis), the Variance Concentration block, and the Risk Register Coverage block.

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. 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

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

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

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.

12.6 Variance Share on the Tornado

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.

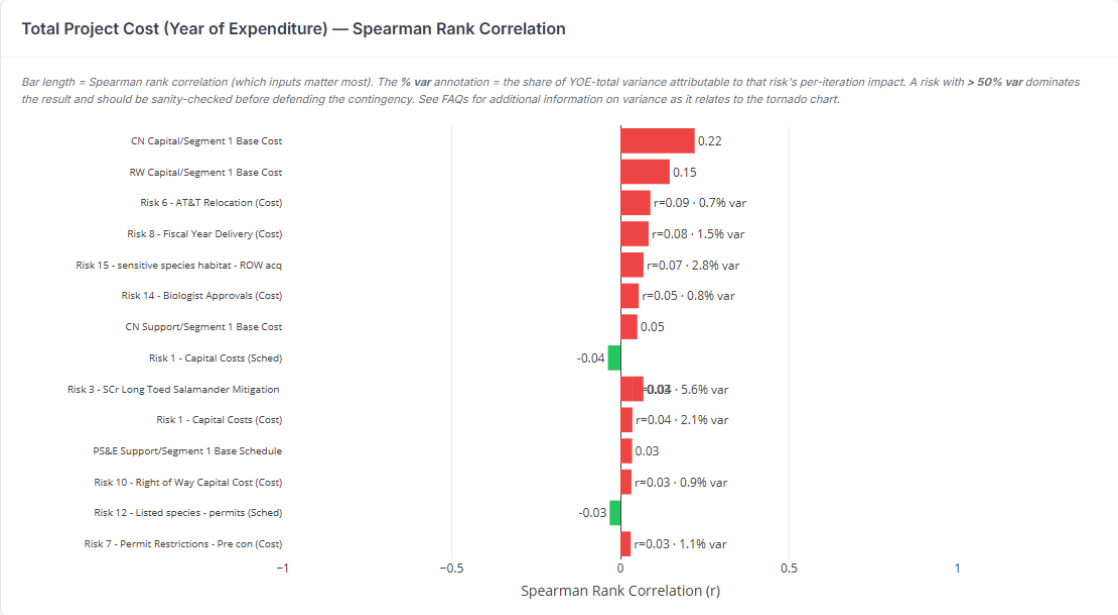

Figure 12.3: YOE Spearman tornado showing the new "% var" annotation on Risk Register rows. Base-cost rows and schedule-axis entries show only the Spearman r value (see why below).

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.

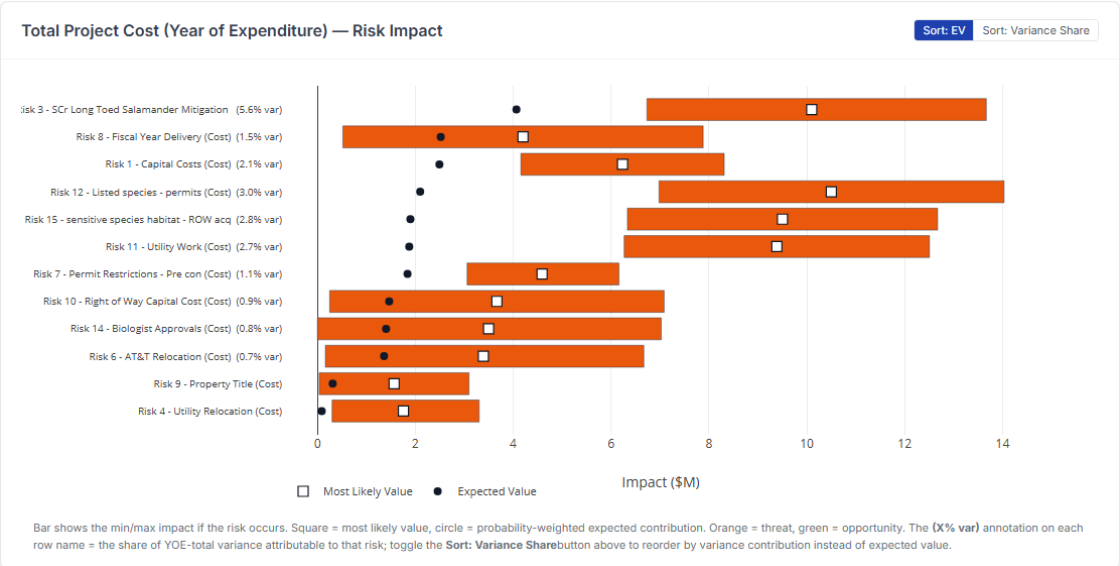

Figure 12.4: Risk Impact tornado with "(X.X% var)" annotations and the Sort: EV / Sort: Variance Share toggle in the card header.

12.7 What-If Sensitivity Tab

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.

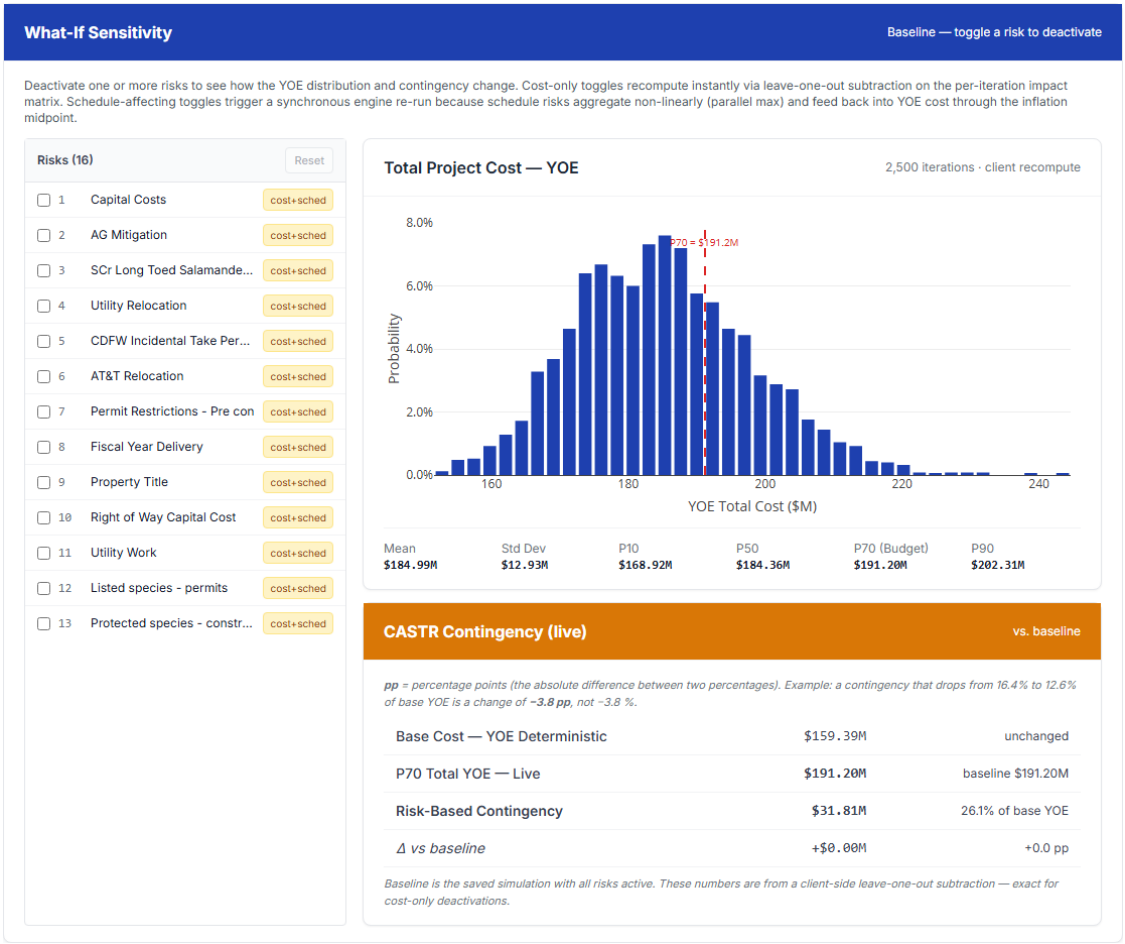

Figure 12.5: What-If Sensitivity tab in its baseline state, with all 16 risks active.

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.

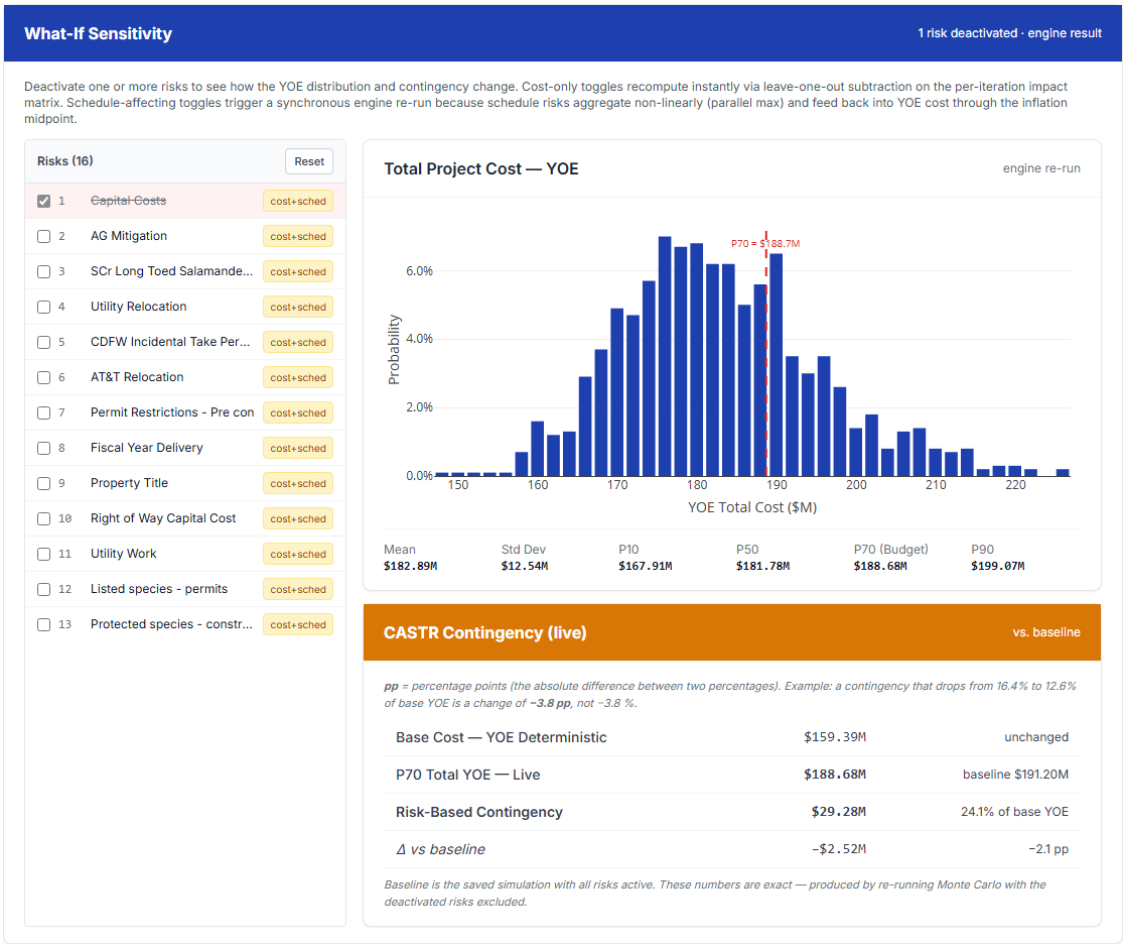

Figure 12.6: What-If Sensitivity tab after deactivating Risk #1 (Capital Costs, a cost+sched risk). The engine re-run badge under the chart title shows the path; the contingency strip's Δ vs baseline row reads −2.52M / −2.1 pp.

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.

12.7.5 Duration and Completion Date views

A Metric dropdown sits above the chart with three options: Cost (YOE), Duration, and Completion Date. The selected metric drives both the histogram and the live-vs-baseline strip below it.

Cost (YOE). The default view. Histogram of YOE cost across iterations with a P-budget marker; the dollar contingency strip shows baseline vs. live numbers.

Duration. Histogram of Total Project Duration in months (calendar time). Driven by the engine result for the current toggle state, or the baseline result when no schedule-affecting risk is deactivated. A Schedule Delta (live) strip appears below the chart, showing the baseline Duration P-budget alongside the live Duration P-budget so you can read the swing without leaving the panel.

Completion Date. Histogram of Total Project Completion Date across iterations. The Schedule Delta (live) strip carries over from the Duration view but reads in calendar dates rather than months. Requires base-uncertainty rows on the project for a meaningful spread; otherwise the chart collapses to a single bin at the planned end date.

Dollar contingency strip stays visible. Switching the Metric dropdown does NOT hide the existing CASTR Contingency (live) card — the dollar strip always reflects the current toggle state regardless of which metric is plotted above. This makes it easy to flip between Cost / Duration / Completion Date views while keeping the dollar headline in sight.

Per-axis Threat / Opportunity chips. Each axis on the risk checklist gets its own chip, colored by polarity: red for Threat, green for Opportunity, gray when the polarity is unknown. A risk with both axes populated shows two chips, so a mixed-polarity risk reads at a glance — for example a red Cost chip alongside a green Schedule chip on the market-conditions example from §5.1.7.

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.
Response (formerly Mitigation)	An intervention that reduces a risk's probability of occurrence and/or its cost or schedule impact. Entered per-risk through the Response Details panel on Risk Identification. CASTR tracks the pre-response (original) and post-response (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-response cost or schedule distribution for a risk with a planned response, representing the risk's expected behaviour if the response strategy is implemented. Edited through the Risk Details panel alongside the pre-response 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
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: 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” 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

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.

“Appendix D: Tips, Tricks, and Best Practices”

This appendix collects the practical tips, shortcuts, and habits that long-time CASTR users rely on to get clean numbers out of the tool quickly and to defend those numbers in front of executives. Each section maps loosely to the manual chapter that covers the underlying feature.

D.1 Setup and Navigation

Demo project first. New installs ship with the *Rural Highway Widening (4 to 6 Lanes) - Demo* project. Open and run it end-to-end before touching real data — every feature (PoS@P70, tornado, MI/ME dependencies, inflation sensitivity, response ROI) is wired up so you can see what “good” looks like.

Trust the Progress Checklist. When a project is open, the sidebar checklist is the canonical setup order (Estimate Last Update Date → Segments/Phases → Inflation Tables → Base Variability → Risks → Simulation). Click an unchecked item to jump straight to it instead of hunting through the navigation.

Use the Guide and FAQ panels. Every page has both. The Guide is a top-to-bottom walkthrough of the page; the FAQ answers targeted “how do I…” or “why is this happening?” questions. They auto-close each other so the workspace stays clean.

Click the question-mark help icons. They are scattered next to field labels and table headers. Each opens a popover *and* deep-links into the manual section that explains the math.

Set the Estimate Last Update Date early. Inflation is computed from this date forward, so getting it wrong silently shifts every YOE number on the project.

D.2 Data Entry and Shortcuts

Paste columns directly from Excel. Click the first cell in the target column and press Ctrl+V — CASTR spreads the column down for you. This works in Base Variability and most data-entry tables.

Use Enter to add a new risk row. In the Risk Identification Triage View, place the cursor in the risk-event name field and press Enter. Faster than clicking the Add Risk button when you are brainstorming with a workshop group.

Use the Starter Risk Library. Instead of typing every risk from scratch, the Starter Library picker inserts common predefined risks. Edit the wording afterward to match the project-specific language.

Do not include prior or fixed costs in Cost Est, CY. Enter them in the dedicated fields above the Base Variability table — they are added to the project total but are not inflated and not subject to probabilistic sampling.

Keep Base Cost and Duration variability small (≤10%). This represents *inherent* estimating variability, not risk events. If you are tempted to set it to 25%, you probably want a Risk Register entry instead.

Skip Market Conditions for PE phases. PE costs are staff and consultant hours, not market-bid construction. The bundled demo omits PE from market-condition rows; mirror that pattern on your own projects.

Watch for red Market Conditions cells. Red highlighting means a gap exists between BTP-High and AP-Low (or AP-High and WTP-Low). Hover the cell for the fix. Gaps create zero-density regions in the multiplier distribution and distort the simulation.

Read the Market Conditions Scenario Distribution chart before running. On the Market Conditions page, just below the data table, the Scenario Distribution chart visualizes the BTP / AP / WTP mixture for whichever phase-segment you select. If the solid mixture line peaks well off 1.00 you are biasing cost up or down — usually unintentionally. Also watch for gaps between the dashed curves and check that the dashed peak heights actually reflect your probability allocation.

Snap per-phase Project Wide dates to the phase span. With two or more user segments, Base Variability shows a "PE / Project Wide" (and RW/PW, CN/PW) row alongside the per-segment rows. The shield-S **Snap** button at the right end of each PW row fills its Start and End dates to span all sibling rows in that phase — earliest start to latest end. The button is highlighted in amber with a "Snap" label while the PW row's dates are blank, then quiets to a plain icon once the dates are aligned. Hover the icon for a preview of the dates it will write. Snap is one-click; if you later adjust PE/Seg 1 or PE/Seg 2 dates, click Snap again to re-align (CASTR won't silently overwrite PW dates you may have intentionally narrowed).

D.3 Risk Identification and Register

Use the All / Res / Non-Res filter chips to triage what still needs a response plan versus what has already been handled.

Combine Heat Map cells to filter the Triage table. Click multiple cells in the Cost and Schedule Heat Maps in turn — clicks combine, so you can isolate clusters of risks (for example, high-cost AND high-schedule) instead of scrolling the entire table.

Use the Contingency Envelope "State" column to steer the workshop. The State column tells you which phase-segments most need attention, so the team can spend its time where it matters instead of reviewing every risk equally.

The Contingency Envelope is a sanity check, not a hard rule. It provides a reasonableness range based on the project's design-completion stage. Treat it as a check on your overall reasoning, not a target you have to hit.

Treat seeded distributions as drafts. When you generate risk distributions from qualitative 1–5 triage scores via the Seed button, the values are starting points only.

Seeding does NOT lock the risk — open the Distribution Editor at any point to refine the low/mode/high values or change the distribution type entirely.

Sync to Risk Register early and often. Reg # always equals the Risk Identification row number, so re-syncing is safe and idempotent. Response fields carry forward, and RES is protected from being overwritten once you have set it.

Use the green-shield icon as your visual sweep. On Risk Identification, mitigated risks turn the Details button green with a shield icon. Scan the column to see what is planned at a glance.

Drag the Risk Details dialog out of the way. The dialog is draggable from its title bar — useful for cross-checking the underlying register row while editing residuals.

Pre-populate residuals from the original. The residual distribution editor seeds from the original distribution the first time you open it. Adjust only what changes — do not rebuild from scratch.

Use the pre vs post Preview overlay. Red is the original (pre-response) curve; green is the residual (post-response) curve, drawn on the same axes. Calibrate residual parameters by comparison rather than guessing.

D.4 The Four-Step Response Workflow

Plan ≠ Active. Setting Risk Response = Mitigate captures the *plan*. The simulation does NOT use it until you tick the **RES** checkbox on the Risk Register.

Compare-Both wastes time if no RES is checked. When every RES is unchecked, Post-Response and Pre-Response produce identical numbers. Switch run modes only after RES is committed on the risks whose responses are funded.

Use the Response What-If panel for live experiments. Toggles in the panel are *local* — they never write back to the register, so you can explore "what if we drop response #7?" without losing your committed plan.

D.5 Simulation

Start with default settings. Pre-Response, default iteration count, P70 budget percentile. Verify your inputs produce sensible numbers before tuning.

Tailor the Early-Stop policy to your audience.

- *Drafts and live workshop runs* — leave Early-Stop ON to generate defensible headline numbers quickly.
- *Formal contingency memos, reports, or projects with extreme tail-buster risks* — turn Early-Stop OFF and manually run a higher iteration count (typically 20,000–50,000) for perfect reproducibility and stable tail-driver rankings.

Watch the tail-driver confidence chips. Green = Reliable; rankings stable across re-runs. Amber = Review; positions 3+ may shift on a re-run. Red = Insufficient; treat the rankings as indicative only, and re-run with more iterations before defending the result.

Soften dense correlations. If you link too many risks together in the Correlation Editor, you may accidentally create a mathematically invalid (non-PSD) matrix that the simulator will auto-correct, sometimes with large adjustments that depart from your model intent. Reduce coefficients (r) into the 0.4–0.6 range to keep the matrix valid.

Match your method to your audience. PoS (Percentile of Sum) is defensible and standard. SoSP and SoPSP are sums-of-percentiles — useful for some agency rules but the project-total chart shows a *reconstructed* distribution because the headline budget is no longer a single point on the raw distribution.

D.6 Reading Results

Always compare Mean vs. Median. If the Mean is materially higher than the Median (P50), you have meaningful tail risk and should defend at P70 or higher rather than P50.

Use the Contribution to Budget chart for executives. A step-by-step waterfall (Base Cost CY → + Base Inflation YOE → + Contingency YOE, or three contingency sub-steps when Component Analysis has been run) that sums exactly to the headline budget. The cumulative buildup tells leadership not just *what* the budget contains but *how* the team arrived at it — the most defensible visual for a finance committee.

Marginal P70 cost AND P70 date is NOT joint P70. The Joint Confidence tab makes this explicit. Drag both crosshairs outward until joint coverage hits your target — there is no single point that hits joint P70, only a contour curve, and you choose the trade-off you can defend.

Use both tornado modes. Spearman is the classical sensitivity view — "which risks drive variability?" — and is the right lens for chasing tail risk. Risk Impact ranks risks by contribution to mean cost in dollars — "which risks add the most expected dollars to the headline?" — and is the right lens for executive briefings. Many CASTR users include both views in their reports.

Toggle YOE vs. CY on the tornado. Inflation can change which risks dominate, especially on long-duration projects.

D.7 Inflation Sensitivity

One inflation table is fine if the rate is uniform. If your project applies the exact same inflation rate across all phases, you do not need multiple tables — create one inflation table and mark it Project-Wide.

Use 2–3 anchors for the Sensitivity Curve. Three anchors yield an exact quadratic; two yield linear; one yields a flat reference. More anchors equal a more honest curve outside the bracket.

Do not defend extrapolated values. The Sensitivity Curve auto-extends roughly 50% beyond your scenario bracket, but the chart caption flags it as directional only.

Schedule charts are scenario-invariant. They render identically across scenarios — a confirmation, not a bug.

D.8 Reporting and Exports

Word export is your one-click defense package. Project Information, summary stats, Contingency table, Response ROI, and presentation-sized chart images all in one .docx. Tick "Include all inflation scenarios" when scenarios are configured and you want a side-by-side defense.

Toasts pause on hover. If a save or import toast disappears too quickly, hover it next time — the dismissal timer pauses while your cursor is over it.

D.9 Data Hygiene, Recovery, and Licensing

Always back up before deleting a project. Deleted projects cannot be recovered in CASTR. Export a backup archive from the Backup / Restore menu before permanently trashing a project.

The recovery card never deletes data. *Try again* first; *Back to projects* second; *Reload page* only as a last resort. Screenshot the card before reloading if support might need to look at it.

License is machine-locked but rehostable. Up to two free rehosts per rolling 12 months — email michael.smith935@gmail.com with confirmation that the old machine is decommissioned.

A lapsed Upgrade Coverage Window is fine. Your license keeps working forever on the version you have. Renew only if you want releases published *after* your window expires.

D.10 Workflow Rhythm

Iterate small, save often. CASTR auto-saves every row, but the rhythm that works in workshops is to enter a few rows, glance at the Scenario Distribution chart or the Contingency Envelope, adjust, and keep going. You will catch input mistakes immediately instead of at the end.

Re-run after model changes. Old runs in the database keep their original results JSON; new metrics (calendar-time duration, Risk Register columns, joint plot data) populate only on fresh runs. If you see "—" or a "Joint plot unavailable" placeholder, just re-run the simulation.

Keep one “clean baseline” run per project. Before you start what-if’ing responses, save a Pre-Response run as your reference. Comparison only means something when there is a fixed baseline to compare against.