

CVI Driver Explorer: Overview

DARING, The Dunne Advisory Risk Intelligence Group

Bottom line: This is an example of how a complicated problem can be approached by embracing its complexity through data and methodology, rather than avoiding it. The result is a set of strategies that make the problem more tractable, and pathways toward risk management, resiliency, and opportunities for growth.

What This Is

An interactive 3D network that turns the U.S. Climate Vulnerability Index (Lewis et al. 2023) from a static score into an explorable structure. It is built from the CVI's own published tract-level data (73,057 census tracts and 184 indicators, rolled into 37 subcategories across 7 categories) with two layers on top:

- **Correlation network:** each subcategory is a node; edges are Pearson correlations between subcategories' tract-level vulnerability scores ($|r| \geq 0.30$). This reveals which parts of the index move together (or trade off against each other) in the real data, not just how the index's own hierarchy is organized.
- **Severity-band explorer:** a slider steps through the overall CVI score in 10-point bands (1–10 through 90–100). At each step, the network resolves to show the top 10 subcategories most elevated in that band. Everything else recedes to grey but stays visible, so the viewer sees simplicity emerge from complexity rather than a fixed answer.

The tool opens on the full, tangled picture (every node, every correlation) by design. That opening state is the argument: the problem looks impossibly complex until you have a way to move through it.

What It Does Not Do

- **It does not establish causation.** An edge between two subcategories means they co-vary across census tracts, not that one causes the other. High betweenness or a strong correlation identifies a subcategory worth investigating as a possible leverage point, not a proven one.
- **It does not model intervention effects.** There is no mechanism for simulating “if we reduce X by 10%, Y falls by Z%.” That would require a causal model (e.g., a literature- or expert-elicited DAG), a materially larger undertaking than this tool.
- **It does not weight by magnitude of harm or population.** A subcategory can rank high in the network while affecting relatively few people, or rank lower while affecting many. The tool is blind to that unless population weighting is added as a separate layer.
- **It carries the CVI's own assumptions forward,** most notably equal weighting of every indicator within a subcategory and every subcategory within a category. Reweighting would require rebuilding the underlying percentile calculations, not just adjusting the visualization.

- **It is ecological, not individual.** All relationships describe census tracts (places), not households or people. Two subcategories correlating across tracts does not mean the same residents experience both.

How It Can Be Used to Supplement the CVI

The published CVI gives a single score per census tract. This tool sits alongside it as a diagnostic layer that answers a different question: *not* “how vulnerable is this place,” but “what is actually elevated here, and does that change as vulnerability worsens?”

- **Reveals a severity handover.** Moving through the bands shows the profile of what's elevated is not static: at moderate vulnerability (50th–60th percentile), environmental exposure (pollution, land use) dominates. By the most severe tail (90th–100th percentile), health outcomes, infrastructure (communications), and disaster mortality take over. A single “top 5 drivers” list would have hidden that shift.
- **Surfaces genuine trade-offs the raw score can't show.** Some subcategories move in opposite directions across tracts (e.g., baseline environmental burden and extreme-weather risk correlate negatively), meaning the places carrying the worst legacy pollution burden are often not the same places facing the sharpest rise in extreme weather. That distinction matters for where to target different kinds of interventions.
- **Turns a 184-indicator, 37-component index into something a non-technical audience can navigate.** The severity slider is built specifically to take a viewer from “this looks impossibly complex” to “here are the 10 things that matter most at this level of severity,” in real time.

How Your Organization Can Use It

- **Client-facing resilience strategy.** Walk a client or partner live through the slider to build a shared, real-time understanding of where their community or portfolio sits on the vulnerability spectrum, and which levers matter at that specific point, rather than a generic checklist.
- **Prioritization exercises.** Use the per-band top-10 lists as a starting point for structured discussion: which of the elevated drivers are actually addressable by this organization, versus which require partners or upstream policy change.
- **Board and stakeholder communication.** The “full complexity → resolved simplicity” interaction is built to make an abstract, technical problem legible in a meeting: open on the tangle, then walk the slider live to make the case for a phased strategy without needing the audience to read a methods section first.
- **A template for other complex-systems problems.** The same approach (correlate real data instead of assuming a hierarchy, threshold to cut noise, and let a severity dimension reveal shifting priorities) could be reapplied to other DARING/IRIS engagements involving many interacting risk factors, not just climate vulnerability.