🌻 Combining links into bundles

In most projects, the data contains many repeated causal claims with the **same cause and the same effect** (often across many sources). We call these **bundles** (or **co-terminal link bundles**).

This extension is about:

- turning many individual coded claims into a smaller set of **bundle rows**, and
- being clear about what we mean when we say "a link" on a map.

Why bundling is useful (for practitioners)

- **Simpler maps**: instead of 200 separate "A → B" rows, you see one "A → B" bundle with counts.
- **Clearer evidence signals**: you can read "how widely shared" (sources) vs "how often said" (citations).
- **Better reporting**: it's easy to state things like "7 sources mentioned A → B".

What you get (in plain terms)

Instead of one row per coded claim, you get one row per **unique cause→effect pairing** (after whatever label transforms you applied).

Each bundle row can show:

- a readable key like `cause >> effect`
- **source count** (how many distinct sources made at least one claim of this form)
- **citation count** (how many coded claims / rows are in the bundle)
- optional summaries like mean sentiment (if you use sentiment)

How to read a "link" on a map (interpretation rule)

On maps and in tables we often still say "link", but it usually means:

> a **bundle** representing "many similar claims that cause influences effect".

So if a link label says "7 sources / 12 citations", read it as:

- 12 coded claims were bundled together, coming from
- 7 distinct sources.

Practical cautions

- **Bundling happens after transforms**: if you zoom/collapse/combine opposites/cluster first, you are bundling the *transformed labels*, not the raw labels. That's often what you want, but be deliberate.
- **Counts are evidence volume, not effect size**: a frequent bundle means "often claimed", not "strong causal effect".

Formal notes (optional)

The filter operates on the current links table (one row per coded claim / citation) and produces a derived table with fewer rows by grouping on the current (possibly transformed) labels.

- Bundle key: (cause label, effect label)
- One output row = one bundle (one unique cause→effect pair)

The bundled output adds aggregate columns such as:

- `bundle`: a readable key like `cause >> effect`
- `citation_count`: number of underlying link rows in the bundle
- `source_count`: number of distinct sources contributing at least one link row to the bundle

Optional further summaries can be computed from the underlying rows (e.g. `mean_sentiment`, per-tag counts, per-group counts).