🌻 Working Papers.

Chapter Contents.

Working papers hub (project overview)

> This folder is a set of working papers about **causal mapping as qualitative evidence management**: we code *reported causal influence claims* from text as a links table with provenance, then analyse the resulting evidence base through explicit transforms and queries.

The overall aim is to keep the core representation **minimal and auditable**, while still supporting powerful downstream analysis (filter pipelines, standardisation/recoding, coverage/fit diagnostics), including workflows that use LLMs as low-level assistants for extraction and labelling.

Core papers (start here)

- Minimalist coding for causal mapping: the core coding stance ("barefoot" link coding), why it is useful, and where it breaks.
- A formalisation of causal mapping: companion spec—data structures + conservative rules for aggregation/query.
- Causal mapping as causal QDA: positioning for qualitative methods / CAQDAS audiences.

Practical extensions (operations on a links table)

- Magnetisation: soft recoding with "magnets" (standardise labels at scale without re-coding quotes).

- A simple measure of the goodness of fit of a causal theory to a text corpus: coverage-style diagnostics for ToC fit.
- Combining opposites, sentiment and despite-claims: opposites transforms, sentiment as an annotation layer, and "despite" link typing.
- Hierarchical coding: hierarchical labels (`;`) and zoom-style simplification.

Related notes / fragments / examples

- !!!Qualitative Split-Apply-Combine: small-Q framing; causal mapping as a SAC variant; where genAI fits.
- 250! causal mapping turns QDA on its head: a short argument/fragment (kept for reuse).
- Conversational AI -- Analysing Central Bank speeches: worked example of "clerk vs architect" (auto-extraction + magnet-style structuring).

structured, auditable qualitative model (a network of causal claims) that can be queried using a transparent library of operations (filtering, path tracing, label transforms, bundling). This gives researchers a way to keep qualitative judgement central while making key intermediate products more reproducible, checkable, and scalable—especially when using AI as a constrained, low-level coding assistant rather than a black-box analyst.

A simple measure of the goodness of fit of a causal theory to a text corpus

Suppose an evaluation team has a corpus of interviews and progress reports, plus (at least) two candidate theories of change (ToCs): an original one and a revised one. A practical question is: **which ToC better fits the narrative evidence**?

Magnetisation

After inductive causal coding (manual or AI-assisted), you typically end up with **many overlapping factor labels** (“rising prices”, “inflation”, “cost of living increases”, …). *Magnetisation* (aka **soft recoding**) is a fast, transparent way to standardise these labels **without re-coding the original text**: you supply a list of target labels (“magnets”), and each existing label is reassigned to its closest magnet by semantic similarity (using embeddings). Unmatched labels can be kept (default) or dropped.

Conversational AI – Analysing Central Bank speeches

Abstract

Hierarchical coding

See also: [[000 Minimalist coding for causal mapping]]; [[015 Combining opposites, sentiment and despite-claims]]; [[005 A formalisation of causal mapping]].

Rigour and causal pathways

Thanks guys, there's lots to like here and lots to agree with. Helping to find what causal hypotheses to focus on… What do you have to say about *pathways* as opposed to individual links/mechanisms? If we called this approach "quality and rigour in causal mechanisms evaluation" would that miss anything?
Disclaimer: for me, the logic around how links might combine into pathways and what that means for evaluation, that's the most exciting part. e.g. how might this intervention influence an outcome which might be multiple steps downstream of it?

