Improving Outcome Harvesting maps with source tracing

31 Oct 2026

For outcome harvesters, and for evaluators who commission or review a harvest. You do not need to know anything about causal mapping to follow it.

Work in progress. We draw on causal mapping's published method and on one pilot that mapped a harvest's outcome descriptions. We have not yet applied source tracing to a harvest's own outcome map.

In short

- **Many harvests end with an outcome map**: the outcomes drawn as boxes, with an arrow wherever one outcome helped to bring about another. It is usually informal.
- **An outcome map is a causal map**, so it inherits causal mapping's main hazard: a chain of arrows can read as a story that nobody told.
- **The hazard is the transitivity trap.** If one account says A led to B, and a different account says B led to C, it does not follow that A led to C.
- **Source tracing is the remedy.** Show only the chains that a single source tells from end to end.
- **A chain that survives source tracing carries the warrant of one account**, and that is what lets a harvester reason along chains of more than one link. A chain that does not survive is a composite: a hypothesis to put to substantiation, not a finding.
- **It can be done by hand**, with a source written against every arrow, or in software that does source tracing, such as Causal Map.

See also: Causal mapping can complement Outcome Harvesting; Outcome harvesting; The transitivity trap.

Outcome maps are causal maps

In outcome harvesting, an outcome is a change in what a social actor does: their behaviour, relationships, activities, policies or practices. Each outcome description says who changed, what changed, why it matters, and how the intervention contributed (Wilson-Grau 2018). The contribution part is already a causal claim, and it comes from a named source: the change agent who described it, or the substantiator who confirmed it.

Harvesters have long found it hard to say how the outcomes in a harvest relate to one another. Some count outcomes by type of actor and chart them; others lay them out on a timeline to work out which made which possible (Britt et al. 2025). Many go one step further and draw an outcome map, where an arrow from one outcome to another means the first helped to bring about the second.

That map is a causal map: a set of causal claims, each of which came from somewhere, drawn as a network. Causal mapping has been used for decades across many disciplines, and it has rules for reading such a network that harvesters can borrow. The most useful is to treat the map as a store of evidence rather than as a model of what happened. In our account of causal mapping for evaluators, causal maps are "repositories of evidence for those pathways" (Powell et al. 2024).

Where the arrows come from

An arrow in an outcome map can rest on quite different things:

- a change agent's own account, in an outcome description, that this outcome followed from that one;
- a substantiator's statement;
- something said in an interview that never became an outcome description;
- the harvester's own inference, often from the order in which things happened.

Informal maps rarely record which. For a single arrow this matters little, since a reader can go back to the description and check. It matters a great deal as soon as somebody reads along a chain, which is what an outcome map invites them to do.

The transitivity trap

Our standard example comes from farmers (Powell et al. 2024). A pig farmer says: "I received cash grant compensation for pig diseases, so I had more cash." A wheat farmer says: "I had more cash, so I bought more seeds." Put the two on one map and a chain appears: pig-disease grants lead to more cash, which leads to more seed. But the first link only makes sense for pig farmers and the second only for wheat farmers. In general, from A→B in one context and B→C in another, we can conclude A→C only where the two contexts overlap, and a map does not tell us whether they do.

The same thing happens in a harvest. Take an invented example:

- A head teacher's outcome description says that after the project's training, the district education officer allowed schools to broadcast lessons on community radio.
- A community leader in a different district says that once lessons were on the radio, parents in her village began to supervise their children's home learning.

On the outcome map these become training → district permission → radio lessons → parents supervise home learning. It reads as the project's contribution to what parents did. But the permission was for one district and the parents were in another, where the radio lessons may have come from somewhere else entirely. Each arrow is sound. The chain is not.

Nothing on an ordinary outcome map distinguishes this chain from one that a single change agent described from start to finish. Both look the same, and the longer the chain, the more joins it has where this can happen.

Source tracing

Source tracing keeps only the chains that can be found within a single source: one person's account, told link by link from the first factor to the last. If source P says A→B and source W says B→C, the chain A→B→C does not appear, because neither source told it. If source Q says A→B and B→C, it does.

In a harvest, a source is whatever you would go back to when checking an arrow: an outcome description, an interview with a change agent, or a substantiator's statement. Which unit you choose is a judgement. An outcome description is the natural unit, because it is the account the method already treats as the record.

Source tracing changes what a harvester can say:

- **Longer chains get a warrant.** A chain of three or four links that survives source tracing is somebody's account of a mechanism, not three fragments placed side by side. The harvester can report it as such, and say whose account it is.
- **Chains can be counted.** Instead of counting how many sources mention each arrow, count how many tell the whole chain. "Five change agents described the route from the training to parents' supervision" is a far stronger contribution claim than five separate arrows each mentioned by somebody.
- **The composite chains become the substantiation list.** Every join that no single source told is where the map is weakest. Put those joins to substantiators first.
- **Gaps become visible.** Trace from the intervention to each domain-level outcome and list which have at least one complete chain behind them and which have none. An outcome the programme expected to influence, with no complete account linking the two, is worth the sense-making workshop's attention.

What it costs

Source tracing is conservative, and it has three costs that a harvester should state.

- **It hides chains that may be real.** Part of the point of causal mapping is the surprise of fitting together fragments that nobody connected, as an evaluator does with evidence from different phases of a programme. A traced map leaves those out. Keep both maps, and say which one a reader is looking at: the composite map for generating hypotheses, the traced map for claims.
- **One account is still one account.** A chain that survives source tracing has the warrant of the person who told it, no more. It still needs substantiation, and even a single source's chain is worth checking with them: did A really lead to B, which led to C, as part of one process?
- **Labels decide what counts as the same outcome.** Two descriptions that both mention "community support for learning" may not mean the same thing. Whether to give them one label is the identity question, which comes before the transitivity trap. Trace first, then merge labels for presentation, because merging first can create chains that were never in the data.

Doing it

By hand, the essentials are few:

- Write a source against every arrow, even when the map is a sketch on a flipchart.
- Draw arrows within one account as solid lines and joins between accounts as dashed lines.
- Before writing a contribution claim that runs along a chain, check whether one source told the whole chain. If not, call it a hypothesis and put the dashed join to a substantiator.
- In the harvest itself, when a change agent describes an outcome, ask what it led to, and what that led to in turn. Chains told in one conversation are worth more than chains assembled afterwards.

In the Causal Map app, path tracing with source tracing turned on does this for any start and end factors, and keeps chains to a maximum length you choose. In interviews, chains of more than about four links are uncommon and hard to interpret (see path tracing). A From × To table then summarises which routes from the intervention to each outcome have complete support from at least one source, and which have none.

The pilot with Heather Britt mapped thirteen outcome descriptions from one domain of a girls' education evaluation and found that the outcomes influenced one another, with some unexpected contributors such as Ministry officials (Britt et al. 2025). Source tracing is the next step for maps like that one: it tells the reader which of those pathways somebody actually described.

I sometimes use AI to help draft these pages, but the final text is ours.

References

Britt, Powell, & Caldas Cabral (2025). *Strengthening Outcome Harvesting with AI-Assisted Causal Mapping*.

Powell, Copestake, & Remnant (2024). *Causal Mapping for Evaluators*. Evaluation, *30*, 100--119. https://doi.org/10.1177/13563890231196601.

Wilson-Grau (2018). *Outcome Harvesting: Principles, Steps, and Evaluation Applications*. IAP.