

CHAPTER

Causal mapping -- overview

CONTENTS

[Intro](#)

[Causal mapping for outsiders](#)

[A causal map consists of multiple links where a link from X to Y means someone believes X influences Y](#)

[Causal mapping helps make sense of many causal claims from many sources](#)

[Causal mapping starts from what people actually say and what they do not say](#)

[Causal mapping has been used for over 50 years in many disciplines](#)

[Do use causal mapping when you have large numbers of claims from multiple sources, and more open research questions](#)

[Do not use causal mapping if you have limited data or want precise models or specific causal links](#)

[Causal mapping approaches differ in application, construction, analysis and how they deal with multiple sources](#)

[Causal mappers believe that humans are good at thinking in terms of causal nuggets](#)

[Causal mappers believe that humans are good at thinking in terms of causal nuggets](#)

[Causal mapping has been used for over 50 years in many disciplines](#)

[Causal mapping differs from related approaches - epistemic, less predictive, unsophisticated, many links, many sources, unclear boundaries](#)

[Task 3 -- Analysing data, Answering questions](#)

[Granularity, generalisability and chunking are coding problems for causal mapping too](#)

[Granularity, generalisability and chunking are coding problems for causal mapping too](#)

[Task 3 -- Analysing data, Answering questions](#)

[Granularity, generalisability and chunking are coding problems for causal mapping too](#)

[Causal mapping is easier if we are realist about causation](#)

[Causal mapping is good at coping with messiness and complexity](#)

[Granularity, generalisability and chunking are coding problems for causal mapping too](#)

Intro

What is causal mapping? What are its strengths and weaknesses?
How does a causal map differ from a systems diagram? This chapter has some answers.

Causal mapping for outsiders

What is Causal Mapping?

Causal mapping is a technique to **visualise what people believe causes what** within a complex system. It creates a "mental map" of the cause-and-effect relationships perceived by an individual or a group.

The process starts with **narratives**—such as interview transcripts, reports, or open-ended survey responses. Causal claims within these texts are systematically identified and structured into a network diagram:

- **Nodes (Boxes)** represent the *factors* or *concepts* (e.g., "Better Training").
- **Links (Arrows)** show the *direction of influence* between them.

The Causal Map App

The specialised **Causal Map app** provides a convenient way to do causal mapping. Users can import interviews or reports and "code" them: highlighting causal claims and adding them to the database. Much of this process can optionally be automated using AI, enabling rigorous analysis of larger datasets.

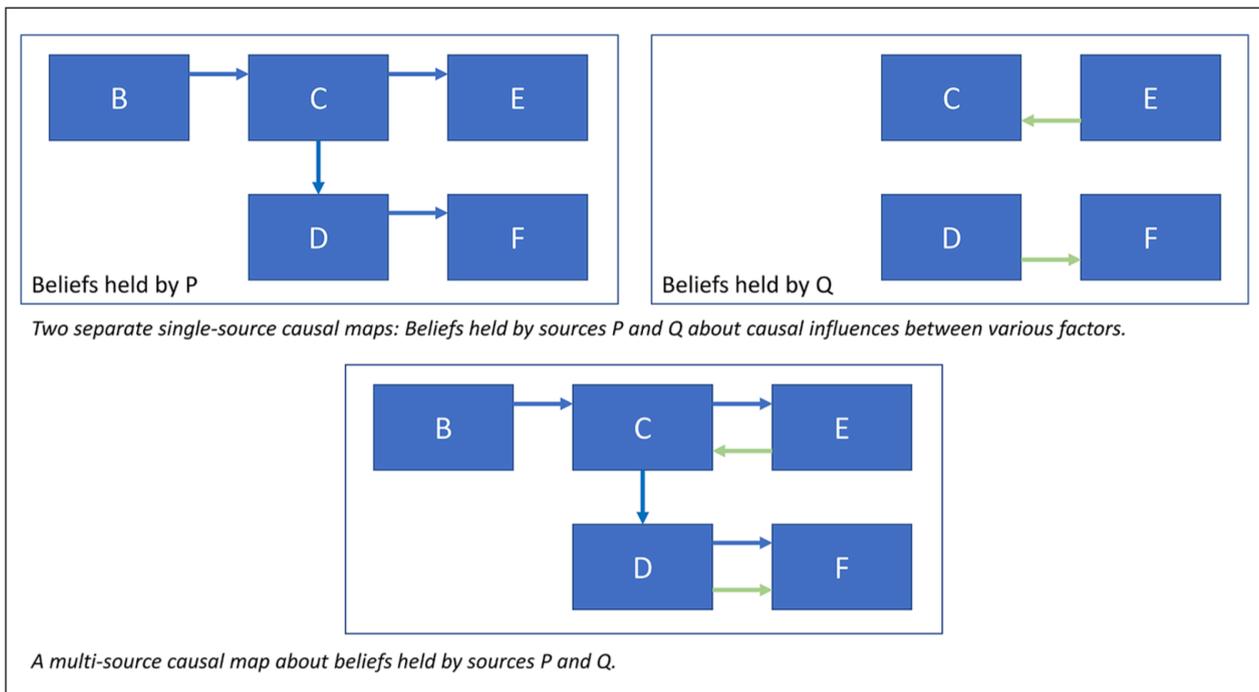
- **Transparency:** Every link in the map is transparently tied back to the **original source quote**. This ensures that outputs are verifiable and avoids acting as a "black box," maintaining the rigour essential for qualitative work.
- **Querying the Map:** The final map is a dynamic model of **causal evidence** that can be actively explored to answer sophisticated questions, such as tracing all direct and indirect links from a single input to a defined outcome.
- **AI as an Assistant:** Generative AI is optionally used as a **tireless, low-level coding assistant** to quickly extract explicit causal claims from text.

A causal map consists of multiple links where a link from X to Y means someone believes X influences Y

Causal mapping helps make sense of many causal claims from many sources

From [Better Evaluation](#)

Causal mapping helps make sense of the causal claims (about "what causes what") that people make in interviews, conversations, and documents. This data is coded, combined, and displayed in the form of maps. These maps show individuals' and groups' mental models and can support further investigation of causal connections.



Top: two separate single-source causal maps: Beliefs held by sources P and Q about causal influences between various factors. Bottom: A combined, multi-source causal map about beliefs held by sources P and Q about causal influences between various factors

Causal mapping is designed for the analysis and visualisation of qualitative data about causal links. It can be used to test an existing theory of change or create collective empirical theories of change about how a program works based on stakeholders' experiences.

People's narratives and reflections about their experiences provide qualitative data that can be coded and displayed as maps to present the cognitive structures ([mental models](#)) of individuals and groups and to support further exploration to understand actual causal connections.

These causal maps can help to answer questions about what people think happened and what they think caused this by building links between different factors, such as different kinds of outcomes

and inputs. Mapping the chains of results and their linkages builds pictures of causal pathways showing the intermediate steps and connections between them.

Causal mapping starts from what people actually say and what they do not say

Causal mapping aims to directly understand and collate the causal claims which people make in narrative (and other) data rather than trying deduce causal connections using statistics or other methods. It starts with what people actually say in real-world contexts and does not rely on heavily pre-structured question formats. Urgent, unexpected, and unwelcome information is treated at face value.

The analyst does not need to have any preconceived conceptual framework; types of causal claims are identified in the data inductively and iteratively. This is a partly creative process, however the decisions made by the analyst are transparent as the underlying text is always available.

At least some of the boundaries of causal mapping research are set by the respondents, not the researchers; what are we going to talk about? What are we not going to talk about?

Causal mapping has been used for over 50 years in many disciplines

From (Powell et al., 2024)

Causal mapping – diagramming beliefs about what causes what – has been used since the 1970s across a range of disciplines from management science to ecology.

The idea of wanting to understand the behaviour of actors in terms of their internal maps of the world can be traced back further to field theory (Tolman, 1948) which influenced Kelly’s ‘personal construct theory’ (Kelly, 1955). A seminal contribution was made by Robert Axelrod in political science, with the book *The Structure of Decision* (Axelrod, 1976). Causal mapping is largely based on ‘concept mapping’ and ‘cognitive mapping’, and sometimes the three terms are used interchangeably, although ‘causal mapping’ strictly involves maps that only include explicit causal links, rather than, for example, relationships like ‘membership’.³ Axelrod’s book presents a comprehensive idiographic approach to how individuals make decisions which he himself mostly refers to as ‘cognitive mapping’ (although his definition makes it clear that all links are causal). An appendix to the book (Wrightson, 1976) gives details about how to code causal links. Bougon et al. (1977) applied a similar approach to a study of the Utrecht Jazz Orchestra as an organisational unit, eliciting ‘cause maps’ from several individual members and amalgamating them. One strand of literature about causal mapping can be located within the wider literature on sensemaking in organisations pioneered by Weick (1995), and applications within organisations were present almost from the start.

By 1990, there were many different applications of similar ideas, including an edited book (Huff, 1990) that offered a unitary approach to ‘concept mapping’ in the United States. Most authors (Ackermann and Alexander, 2016: 892; Clarkson and Hodgkinson, 2005: 319; Fiol and Huff, 1992: 268; Laukkanen, 2012: 2; Narayanan, 2005: 2) use a broadly similar definition of a causal map: A causal map is a diagram, or graphical structure, in which nodes (which we call factors) are joined by directed edges or arrows (which we call links), so that a link from factor C to factor E means that someone (P) believes that C in some sense causally influences E. There is a constructive ambiguity (Eden, 1992) about what a collective map is a map of: While maps constructed as a consensus within a group can plausibly be claimed to map ‘what the group thinks’, this is more problematic for maps constructed post hoc by synthesising individual maps.

We found no significant deviations from this basic definition of a causal map across all the variants of causal mapping reviewed in the following sections, with the caveat that there is variation in how explicit different authors are in describing causal links as representing bare causation as opposed to beliefs about causation.

In the following decades, Eden et al. (1992) applied the approach to understanding and supporting decision-making in organisations, increasingly using the phrase ‘causal mapping’ rather than ‘cognitive mapping’, and they subsequently extended the application of causal maps to fields as varied as risk elicitation and information systems development (Ackermann and Eden,

2011; Ackermann et al., 2014), also developing a series of software packages beginning with Decision Explorer (Ackermann et al., 1996). There is now a wealth of literature on using causal mapping for decision support in organisations (including sophisticated approaches to formalise decision support (Montibeller et al., 2008) and even to rank options (Rodrigues et al., 2017)).

Laukkanen (1994, 2012; Laukkanen and Eriksson, 2013) also wrote extensively on causal mapping and developed a software programme called CMAP3 for processing both idiographic and comparative causal maps by importing, combining and analysing factors and links attributed to one or more sources. A broadly similar approach was taken by Clarkson and Hodgkinson (2005) with their Cognizer approach and software.

References

Powell, Copestake, & Remnant (2024). *Causal Mapping for Evaluators*.
<https://doi.org/10.1177/13563890231196601>.

Do use causal mapping when you have large numbers of claims from multiple sources, and more open research questions

Do not use causal mapping if you have limited data or want precise models or specific causal links

Causal mapping approaches differ in application, construction, analysis and how they deal with multiple sources

Causal mappers believe that humans are good at thinking in terms of causal nuggets

Causal mappers believe that humans are good at thinking in terms of causal nuggets

See also [Causal mapping has been used for over 50 years in many disciplines](#)

[Causal mappers believe that humans are good at thinking in terms of causal nuggets](#)

From [Powell, Copestake, et al. \(2023\)](#)

We claim: our everyday causal understanding is as primary as our perception of, say, colour and arises from more than empirical observations of associations between objects or events; our ability to infer causation goes beyond and is not primarily based on noting correlations. And for all its complexity and intuitive brilliance, it is also just as fallible as our perception of colour or size.

This reaffirms our practice as evaluators of taking the causal claims and opinions of humans (experts and non-experts) seriously (Maxwell, 2004a, 2004b); indeed, this kind of information is the bread and butter of most evaluations.

We can thank [Judea Pearl](#) for promoting the insight that if you want to thrive in this world, you have to understand causality natively. We humans make causal connections [from an early age](#). We wouldn't survive long if we didn't.

GPT-3.5 just about understood causation. GPT-4 and more recent models understand causal connections within text very well.

Our understanding of the world is drenched with causal understanding: information and hypotheses about how things work (mostly accurate enough, sometimes not). It's really hard for us to *not* think causally: the concept of correlation is much harder to understand than the concept of causation.

Causal Inference?

Causal inference is the process of determining whether and how one event or variable brings about another.

Some writers mistakenly assume that only a controlled experiment can "really" provide a route to causal inference.

We would go so far as to say that we don't usually in any conscious cognitive sense *infer* causation -- we just see it, all the time, everywhere. We don't have to teach children to infer causation: we have to teach them to question their perceptions of causation and to distinguish causation from correlation.

Why are humans the best detectors of causation?

1. Evolutionary Adaptation:

Human brains have evolved specifically to detect and act on causal relationships. Survival depends on recognizing which actions lead to which outcomes—e.g., which plants are safe to eat, which animals are dangerous, and how to use tools. This evolutionary pressure has made causal reasoning a core part of human cognition.

2. Intuitive Causal Models:

From infancy, humans build mental models of the world that are fundamentally causal. Children naturally ask "why" questions and seek explanations, not just associations.

3. Generalization and Flexibility:

Humans can generalize causal knowledge across domains. For example, understanding that pushing causes movement can be applied to objects, social situations, and abstract concepts. This flexibility allows humans to detect causation even in novel or ambiguous situations.

4. Counterfactual Reasoning:

Humans often engage in counterfactual thinking—imagining what would happen if things were different. This is a hallmark of causal reasoning and is essential for planning, learning from mistakes, and scientific discovery.

5. Distinguishing Correlation from Causation:

While humans sometimes make errors (e.g., seeing causation where there is only correlation), we are still very good at using context, background knowledge, and intervention to distinguish true causal relationships from coincidence.

6. Social and Cultural Transmission:

Human societies accumulate and transmit causal knowledge across generations through language, stories, and education. This collective causal understanding is a foundation of science, technology, and culture.

7. Observation and Pattern Recognition:

Humans are adept at noticing regularities and anomalies in their environment. We naturally look for patterns—such as temporal precedence (A happens before B), co-occurrence, and changes following interventions—that suggest causal relationships. Even without formal training, people intuitively apply principles like "no effect without a cause" and "causes precede effects."

8. Intervention and Experimentation:

Humans frequently test their causal hypotheses by intervening in the world—changing variables and observing outcomes. This hands-on experimentation, whether in childhood play or scientific research, is a powerful tool for distinguishing causation from mere correlation.

9. Use of Multiple Sources of Evidence:

Humans combine different types of evidence—temporal order, statistical regularities, mechanistic explanations, and observed interventions—to make robust causal inferences. We can weigh conflicting evidence, consider alternative explanations, and update our beliefs as new information arises.

In sum, humans are not just passive recipients of causal information; we are active causal detectives, constantly inferring, testing, and refining our understanding of how the world works. This multifaceted approach to causal inference is what makes us the best detectors of causation.

Conclusion

Causal reasoning is not just a feature of human cognition—it is its backbone. Our ability to detect, infer, and act on causal relationships is what allows us to navigate, survive, and thrive in a complex world. While formalised, controlled experiments are an incredible tool for causal inference in very particular contexts such as some areas of education and economics where multiple very similar causes are regularly followed by multiple very similar effects, if we had to stick to that kind of knowledge and that kind of context we would never be able to get out of bed in the morning, let alone get the kids to school.

Humans remain the best detectors of causation, both individually and collectively.

See also: [400 realist causation](#)

Causal mapping has been used for over 50 years in many disciplines

See also [Causal mapping has been used for over 50 years in many disciplines](#)

From [Powell, Copestake, et al. \(2023\)](#)

Renewed interest in causal mapping may also be reinforced by the ‘causal revolution’ in quantitative data science initiated by Judea Pearl (Pearl, 2000; Pearl and Mackenzie, 2018), which has fundamentally challenged the almost total taboo placed on making or assessing explicit causal claims, which was dominant in statistics for much of the twentieth century (Powell, 2018), and this has in turn helped rekindle interest in explicitly addressing causation using qualitative methods.

Causal mapping differs from related approaches - epistemic, less predictive, unsophisticated, many links, many sources, unclear boundaries

Task 3 -- Analysing data, Answering questions

Different approaches to these three tasks are discussed in turn in the following sections.

- [Task 1 -- Gathering narrative data](#)
- [Task 2 -- Coding causal claims as causal qualitative data analysis](#)
- [Task 3 -- Analysing data, Answering questions](#)

Granularity, generalisability and chunking are coding problems for causal mapping too

From (Powell et al., 2024)

How to collect causal claims from which to draw causal maps?

(There is also a whole chapter about this task: [Task 1 -- Introduction](#))

There are a wide variety of options, including in-depth individual interviews (Ackermann and Eden, 2004), reuse of open-ended questions in structured surveys (Jackson and Trochim, 2002), literature reviews (in which ‘sources’ can be documents rather than individuals) and archival or secondary material within which pre-existing causal claims are already made (Copestake, 2020). Other approaches aim to build consensus by using structured collaborative processes, including Delphi studies and PSM (Penn and Barbrook-Johnson, 2019). Guidelines for causal mapping may include procedures for collecting primary data, with forms of elicitation including back-chaining (‘what influenced what?’) and forward-chaining (what resulted, or could result, from this?)

With primary data collection, we can distinguish between relatively closed and open approaches and whether respondents are forced to choose between pre-selected optional answers or can formulate their own (see Table 2). Interviewers may also be guided by a chaining algorithm; for example, they may be instructed to iteratively ask questions like ‘You mentioned X, please could you tell me what were the main factors that influenced X or led to it happening.’

Table 2. Different approaches within primary data collection for causal mapping, with example questions.

Admissible answers / Scope of questions	Explicit: factors are explicitly identified	Implicit: factors are not explicitly named
Closed: questions with a predetermined focus	Which factors in this list influenced this particular event?	What influenced this particular event?
Open: a freer discussion	Identify the biggest change you experienced in relation to X, and list three factors that influenced it	Tell me what has changed for you in the last x years

References

Powell, Copestake, & Remnant (2024). *Causal Mapping for Evaluators*.

<https://doi.org/10.1177/13563890231196601>.

Granularity, generalisability and chunking are coding problems for causal mapping too

From (Powell et al., 2024)

Some approaches such as that suggested by Markiczy and Goldberg (1995) directly elicit causal links from their sources, perhaps by asking respondents to suggest causal links between a predetermined list of causal factors, and thus, after finishing Task 1, are already in a position to create causal maps.

(See also the dedicated chapters on coding...)

- [Task 2 -- Introduction](#)
- [Tasks 2 & 3 -- Introduction](#)

More explicitly, qualitative approaches are faced with Task 2: encoding causal claims in the form of explicit causal links and factors. This task is similar to ordinary qualitative data analysis (QDA), whether done manually or using tools like NVivo, Dedoose and AtlasTI. However, these tools are designed to capture general concepts, rather than claimed causal links between concepts, which is what we need for causal mapping. QDA for causal mapping also starts with a corpus of narrative data, but it does not create causal links between independent concepts that might already have been coded using ordinary non-causal thematic analyses. Rather, in causal QDA, the primary act of coding is to highlight a specific quote from within a statement and identify the causal claim made by simultaneously identifying a pair of causal factors: an ‘influence factor’ and a ‘consequence factor’.

The causal factors only exist as one or other end of a causal link and have no meaning on their own. Each claim forms a link in the visual representation of the causal map. The Axelrod school had its own coding manual describing how to highlight areas of text expressing causal connections and code them as links between causal factors, originally inspired by evaluative assertion analysis (Osgood et al., 1956).

Manual causal coding of text data, like ordinary thematic coding, requires a considerable investment of time and expertise to do well. We now use natural language processing to at least partially automate this; however, the process is essentially the same, and discussion of this is beyond the scope of the present article.

Where do the labels for the causal factors come from? As with ordinary QDA and thematic analysis (Braun and Clarke, 2006), approaches vary in the extent to which they are purely exploratory or seek to confirm prior theory (Copestake, 2014). Exploratory coding entails trying to identify different causal claims embedded in what people say, creating factor labels inductively and iteratively from the narrative data. Different respondents will not, of course, always use precisely the same phrases, and it is a creative challenge to create and curate this list of causal factors. For example, if Alice says ‘Feeling good about the future is one thing that increases your

wellbeing’, is this element ‘Feeling good about the future’ the same as ‘Being confident about tomorrow’ which Bob mentioned earlier? Should we encode them both as the same thing, and if so, what shall we call it? We might choose ‘Positive view of future’, but how well does this cover both cases? Laukkanen (1994) discusses strategies for finding common vocabularies. As in ordinary QDA, analysts will usually find themselves generating an ever-growing list of factors and will need to continually consider how to consolidate it – sometimes using strategies such as hierarchical coding or ‘nesting’ factors (as discussed in the following section).

The alternative to exploratory coding is confirmatory coding, which employs an agreed code book, derived from a ToC and/or from prior studies. QuIP studies mostly use exploratory coding but sometimes supplement labels with additional codes derived from a project’s ToC, for example, ‘attribution coding’ helps to signify which factors explicitly refer to a specific intervention being evaluated (Copestake et al., 2019b: 257). However, careful sequencing matters here because pre-set codes may frame or bias how the coder sees the data (Copestake et al., 2019a). Again, the positionality of the coder matters just as much when doing causal coding as it does for any other form of qualitative data coding.

Combining Tasks 1 and 2

Tasks 1 and 2 result in a coded data set of causal claims, each of which consists of (at the very least) the labels for a pair of causal factors. Those using a more explicit elicitation approach have been able to skip Task 1.

References

Powell, Copestake, & Remnant (2024). *Causal Mapping for Evaluators*.
<https://doi.org/10.1177/13563890231196601>.

Task 3 -- Analysing data, Answering questions

From (Powell et al., 2024)

The extensive causal mapping literature provides many examples of its use to answer evaluation questions (see [Powell, Copestake, et al., 2023](#), p. 110), for example:

- Getting an overview of respondents' "causal landscape". This can be useful for orientation or for particular tasks like triaging masses of information to identify key outcomes and possible causal pathways when planning an Outcome Harvesting (Wilson-Grau & Britt, 2012) or Process Tracing (Befani & Stedman-Bryce, 2017) project.
- Weighing up evidence about contribution: in particular, tracing back and comparing the possibly multiple contributory causes of an important outcome or consequence (Goertz & Mahoney, 2006), or examining effects of causes.
- Reporting key metrics of the causal network, for example, to reveal which factors are most central in the whole network or to identify feedback loops.
- Asking whether the empirical ToC matches the plan ([Powell, Larquemin, et al., 2023](#), p. 7).
- Making comparisons between groups or across timepoints.

One way to simplify is to derive from the global map several smaller maps that focus on different features of the data. For example, maps may selectively forward-chain the multiple consequences of a single cause – including those activities being evaluated: effects of causes (Goertz and Mahoney, 2006) – or trace back to the multiple contributory causes of an anticipated or highly valued outcome or consequence: causes of effects. A series of simpler causal maps, each selected transparently to address a specific question, generally adds more value to an evaluation than a complicated, if comprehensive, single map that is hard to interpret. The downside of this is that selectivity in what is mapped and is not mapped from a single database opens up the possibility of deliberate bias in selection, including omitting to show negative stories.

!

Sets of individual links with the same influence and consequence factor (co-terminal links) are usually represented bundled together as a single line, often with thickness of the line indicating the number of citations, and/or with a label showing the number of links in the bundle. The map has not fundamentally changed, but the visualisation is much simpler.

Simplification - factor and link frequency ▶

Another way to simplify a global causal map is to produce an overview map showing only the most frequently mentioned factors and/or links. Care should be taken if this leads to omitting potentially important but infrequently mentioned evidence about, for example, an unintended consequence of an intervention.

Simplification - hierarchical zooming ▶

Another common way to simplify is to combine sets of very similar factors into one. For example, if hierarchical coding has been used, it is possible (with caveats) to ‘roll up’ lower-level factors (such as health behaviour; hand washing and health behaviour; boiling water) into their higher-level parents (health behaviour), rerouting links to and from the lower-level factors to the parent (Bana e Costa et al., 1999).

Reporting global and local network statistics ▶

Large causal maps can also be analysed quantitatively, including by tabulating which factors are mentioned most often, identifying which are most centrally connected or calculating indicators of overall map density, such as the ratio of links to factors (Klintwall et al., 2023; Nadkarni and Narayanan, 2005). We are wary of the value of summarising maps in this way, not least because results are highly sensitive to the granularity of coding. For example, although a specific factor such as ‘improved health’ might have been mentioned most often, if two subsidiary factors had been used instead (such as ‘improved child health’ and ‘improved adult health’), these two separate factors would not have scored so highly.

References

Befani, & Stedman-Bryce (2017). *Process Tracing and Bayesian Updating for Impact Evaluation*. <http://dx.doi.org/10.1177/1356389016654584>.

Goertz, & Mahoney (2006). *A Tale of Two Cultures: Qualitative and Quantitative Research in the Social Sciences*. Princeton University Press. [12345](#).

Powell, Copestake, & Remnant (2024). *Causal Mapping for Evaluators*. <https://doi.org/10.1177/13563890231196601>.

Wilson-Grau, & Britt (2012). *Outcome Harvesting*.

Granularity, generalisability and chunking are coding problems for causal mapping too

Never confuse the two.

TODO

[Our approach clearly distinguishes evidence from facts and does not automatically warrant causal inferences](#)

Causal mapping is easier if we are realist about causation

Causal mapping is easier if we are realist about causation. We can say that narrative accounts are full of claims about causal powers, that X had the power to affect Y, and X did exercise that power and Y was affected (perhaps in this particular case in spite of or with the assistance of other things).

Causal realism invites us to say that things **have the causal power** to affect other things.

The weird thing is that most physical and natural scientists think about causation in a realist way, but in the social sciences we tell ourselves not to because it isn't scientific (!).

Causal mapping is good at coping with messiness and complexity

From (Powell et al., 2024)

... recognising head-on the ambiguity of much narrative causal data, particularly when confronted with large bodies of data collected in disparate ways. Evaluators must contend with messiness: imprecise system boundaries, differing specification of claimed causal influences and lack of clear or consistent information about what case or group of cases claims refer to. Causal mapping can contend with all this ambiguity rather than shying away from it. It can make use of messy operational data, treating urgent, unexpected and unstructured information at face value. This is made possible by distinguishing clearly between two analytical steps in evaluation: The first is to gather, understand and assemble causal evidence from different sources (those in a position to have useful evidence about relevant causal links and chains) to construct, compare and contrast the evidence for and against different possible causal pathways. By focusing on this task, causal mapping lays a more reliable foundation for the second, often critical, task of using the assembled data to make judgements about what is in fact really happening. This avoids the confusion and ambiguity that often arises when evaluators seek to address both steps simultaneously by constraining what data are collected to fit a prior view of reality which other stakeholders may or may not share.

References

Powell, Copestake, & Remnant (2024). *Causal Mapping for Evaluators*.

<https://doi.org/10.1177/13563890231196601>.

Granularity, generalisability and chunking are coding problems for causal mapping too