

# Task 1 – Gathering causal mapping data

---

## Contents

---

Task 1 – Introduction

---

You can also create consensus causal maps directly

---

Direct links from intervention to outcome are real links too

---

Causal backchaining

---

Changes and states

---

Sources, cases and context in causal mapping and the Causal Map app

---



# TASK 1 – INTRODUCTION

CHAPTER CONTENTS.

---

📅 22 Oct 2025

## Consensus versus multi-source data collection

At Causal Map we are relatively agnostic about data collection. We are most interested in causal evidence and beliefs derived from different sources.

See also:

- [Task 1 – Gathering narrative data](#)
- [Background on data collection with Qualia](#)

PAGES IN THIS CHAPTER

📄 **You can also create consensus causal maps directly**

---

📄 **Direct links from intervention to outcome are real links too**

---

📄 **Causal backchaining**

---

📄 **Changes and states**

---

📄 **Sources, cases and context in causal mapping and the Causal Map app**

---



# YOU CAN ALSO CREATE CONSENSUS CAUSAL MAPS DIRECTLY

---

📅 18 Nov 2025

## Consensus versus multi-source data collection

Some methods related to causal mapping and even some forms of causal mapping itself are not interested in individuals' different causal perspectives: they are primarily aimed at reaching a consensus, expert map straight off without first recording individual viewpoints and attempting to combine them.

TODO

## Related

- [chapter intro](#)



# DIRECT LINKS FROM INTERVENTION TO OUTCOME ARE REAL LINKS TOO

📅 7 Nov 2025

You have a causal map with lots of links from an intervention to a final outcome. It's a really impressive chain.

Then a stakeholder comes along and draws a direct link from intervention to outcome, saying "these trainings are great, they caused the law to get changed!"

You can also create consensus causal maps directly, and in these cases these cheeky direct links are often ignored or curated out of the way. When coding individual sources, and source X says  $A \rightarrow B \rightarrow C$  and then and source Y says  $A \rightarrow C$  it is very tempting to recode source X as if they were really saying  $A \rightarrow B \rightarrow C$ .

For example when source X says

Thanks to the training, our skills increased and so we were able to manage the outreach work better

... and source Y only says

Thanks to the training we were able to manage the outreach work better

... it is very tempting to add the intermediate factor to source Y's story too *even if they did not actually mention it*. If you're going to do that, you should document when and how this is allowed and why.

## Related

- [chapter intro](#)



# CAUSAL BACKCHAINING

---

📅 15 Nov 2025

We have often processed data gathered specifically for evaluation purposes using "causal back-chaining". But we often process secondary data which was not specifically intended for causal coding.

## Related

- [chapter intro](#)



# CHANGES AND STATES

---

📅 15 Nov 2025

When gathering primary data, the way in which questions are asked influences the meaning of the maps and their links. For example, in the QuIP (Copestake et al. 2019) respondents are asked to identify causes of changes, then causes of the causes and so on. This means that most of the factors are already as *changes in something*, such as ‘an improved harvest’ or ‘reduced hunger’). This has implications for how positive and negative statements are combined, as discussed later.

## Related

- [chapter intro](#)
- 

## References

Copestake, Morsink, & Remnant (2019). *Attributing Development Impact: The Qualitative Impact Protocol Case Book*. March 21, Online.



# SOURCES, CASES AND CONTEXT IN CAUSAL MAPPING AND THE CAUSAL MAP APP

📅 12 Aug 2022

## Aim

In a QDA study which interviews 12 farmers about their own experiences, we can code information from source S like “the new seeds brought me better crops” as information about source S. But what do we do with information like “I think most of us had the same successes” or “the farms in the other side of the valley did not have the same success”, or “the seeds were very productive for me but only in the shadier fields”? These challenges arise especially in causal mapping and are very real in our daily work at Causal Map and Bath SDR.

In qualitative data analysis it can be challenging to distinguish between sources and cases, and to decide how to code information which belongs to a different case or cases than the source itself. In Part 1 I explore only the distinction between sources and cases (which we haven't really made so far at Causal Map) without any reference to *causal* QDA. I admit to complete ignorance of how these things are handled in ordinary non-causal QDA, or in QDA software. I'll look that up later.

In part 2 I want to set out an interesting and I think new idea about how causal mapping can give us a different understanding of context and how it relates to cases and sources. The contents of this essay might seem a bit banal to start with but it gets interesting as it goes along.

At the end you'll find a list of quite a few takeaways – things to think about when planning and analysing a study. Most of them are variants of things we already do or issues we are already aware of, but I hope at least I've managed to address them more systematically.

## Part 1: Sources and cases

### Homogenous or heterogeneous sources

Sometimes our sources do not form a single homogenous set: this is often the situation that we as evaluators are in when alongside a clear set of documents we may also have another set of systematic interviews and also perhaps a meta-evaluation parallel to ours but on a different theme.

## Homogenous sources

In causal mapping and when using the Causal Map app in particular, and indeed in QDA in general, when we have multiple sources (variants 2 and 4), when our sources are homogenous we can provide a priori additional dimensions or meta-data (called “additional data” in Causal Map) to our sources like this ...

Source ID	Type	Year	Commissioning Agency
1	Final evaluation	2010	USAID
2	Mid-term evaluation	2009	National government
3	...	...	...

Table 1

... or this ...

Source ID	Age	Gender	Location
1	20-30	M	District A
2	20-30	F	District B
3	...	...	...

Table 2

This makes it easy to filter and combine our information to answer questions like “did the younger sources give different information from the older sources” etc.

## Heterogenous sources

Classic QuIP studies are a good example of heterogeneous source types. Essentially we have separate tables, each of which applies to only part of the set of sources, like this:

Individual Interviews			
Source ID	Age	Gender	Location
I1	20-30	M	District A
I2	20-30	F	District B
I3	...	...	...

Table 3

Source ID	Gender	Location
F1	All Female	District A
F2	Mixed	District B
F3	...	...

Table 4

It is generally not a problem to simply combine heterogeneous datasets as best we can in this kind of fashion:

Source ID	Type	Age	Gender	Location	Gender (FG)
I1	Individual	20-30	M	District A	N/A
I2	Individual	20-30	F	District B	N/A
...		...	...	...	
F1	Group	N/A	N/A	District A	All Female
F2	Group	N/A	N/A	District B	Mixed
		...	...	...	

Table 5

This means of course that it can be a bit trickier to aggregate and answer our research questions using the source meta-data: we can answer questions like “did sources from the different locations give different or similar information” but it is harder to answer similar questions about age, which is (in this case) not defined for the focus groups.

In Causal Map and other QDA software this metadata can conveniently be provided in the form of a table in advance of the actual coding process.

## How are cases and sources related?

Until now we have not clearly distinguished between sources and cases. While the individual interview part of classic QuIP studies maintains a strict equivalence between sources and cases, in other "QuIP-lite" and more generally Causal Map studies, we often have sets of sources reporting about different cases and kinds of cases.

Causal maps often encode information from a number of sources, each describing their own situation. Here we can say that each source is one case and vice versa. Often, where there is no reason to distinguish, we can talk just about “sources”. But the two ideas are quite distinct: the “source” is where the information comes from, and the “case” is what the information is about.

A “classic” QuIP dataset motivates us to think about this distinction. It consists of one part, the individual interviews, made up of multiple sources, each speaking about their own case, and another part which come from multiple focus groups, each of which can be understood as a different source (albeit of a different type from the first part of the dataset) but where the analyst may be in a quandary as to whether to treat the data from each focus group as “about” individual cases consisting of multiple people or whether to try to treat them as “about” multiple cases (the individual members of the group).

More generally the following situations are possible:

1. **One source, one case:** Only one source (or a single consensus of sources), describing one case, (for example, “the water supply system in this part of the country”). As there is only one case, the very concept of case is not very clear or useful. The subject of the map we draw may be expressed in the plural, such as “patients with borderline disorder”, but this plural expression is not important and could equally be expressed in the singular, for example “a typical patient with borderline disorder”. This kind of map is a special case in causal mapping and is more or less equivalent to most of the kinds of systems maps described by (Barbrook-Johnson & Penn 2022). Examples:

- a) a map is drawn by a single expert,

- b) a map is drawn as a collective or participatory synthesis of the views of contributors, as in Participatory Systems Mapping (Barbrook-Johnson & Penn 2021; Wilkinson et al. 2021).

2. **Many sources, one case:** Many sources give information about just one case, e.g. different water systems experts each give their opinion about how one specific water system works, and the information from the different sources is not immediately amalgamated. As with variant 2, if the cases are themselves individuals or groups it might be necessary to specify whether a causal map encodes the beliefs of the source or the (alleged) beliefs of the case.

3. **One source, many cases:** Just one source gives information about many cases, e.g. a manager gives information about how the members of different organisational units think about the way the organisation functions, or a single global evaluation report gives information about causal relationships within a set of different projects.

4. **Many sources, many cases**

- a) **For each source, one case is defined.** For each source, we still know what case it covers. This means that we can easily add a column to the sources table to define the case.

i. **1-1 relationship:** Most commonly, there are many sources, and each source is one case and vice versa. This is for example what happens with the individual interviews in QuIP. Another variant is when we have a set of evaluation studies (sources), each one about a separate project (cases): there is a 1-1 relationship, but the sources are not the cases.

ii. **1-many relationship:** In a meta-evaluation we may have several sources for one case, e.g. several reports on each of several projects.

b) **Cases not defined by sources:**

i. **Cases are still clearly defined:** The classic example is a portfolio meta-evaluation, where multiple evaluation reports each give information about a different subset of a portfolio of projects. For example in a meta-evaluation of interventions after a disaster, there might be an evaluation of all of the health projects, and another final evaluation covering all of the projects implemented by USAID, and so on.

ii. **Cases are not clearly defined:** This is also quite common in many Bath SDR projects which do not follow a strict QuIP format, in which various sources are reporting about, say different projects. Sometimes the identities of the cases become clear and are even enumerated during the coding (Project X, Project Y etc). Alternatively the identity of individual cases is never specified, but still claims are made about subsets of cases such as “women are more likely to have their lives affected by this problem”. At the extreme, sources make claims about multiple different kinds of cases; for example, sometimes they make claims about “people”, and “young people” but also about “businesses”, and “some small businesses” and so on.

In summary:

	Single	Multiple
Sources		
Cases		
Single	1	2
Multiple	3	4

Table 6

Ordinary systems mapping is probably best identified with 1, but arguably includes 2 and 3 as well. Causal mapping has tried at various times to cover all of these variants.

We should remind ourselves also that the distinction between heterogeneous and homogenous sources described above also cuts across variants 2 and 4.

## Coding caseness in causal mapping

Where we have a 1-many relationship between cases and sources (variant 4a) we don't need to worry about how to start coding caseness (to which case does this information apply?) because we already know what case is covered by looking at the source ID. However, in any of the variants we may *additionally* wish to code caseness, for example an interview source talks not only of themselves but about someone else; and in variant 4b we *have* to code caseness, i.e. **we can't provide information about which case is meant until we actually read each causal claim, so we need to be able to create a new "case" column in the table of codings (in Causal Map this is the Links table) during coding, and subsequently add information in this column for each coding to show what case the coding belongs to.** In variant 4a we may not want to treat information given by one source about other sources/cases as *equivalent in value* to self-reported information, or we may choose (as in QuIP) not to code it at all; but in general we may not want to throw away this kind of information.

Different approaches and software applications will have different ways to add caseness information to a coding. Basically all we need is a way to assign a case to a piece of coding. In Causal Map, our approach involves so-called "Quickfields", with which columns of data can be created on the fly.

For example, let's look at the phrase, as part of a meta-evaluation which hopes to provide a combined evaluation of projects X, Y and Z, from an evaluation of project X: "as for project Y, I know their logistics were badly impacted by the flood". We can code the causal part of the claim as usual but add the quickfield *project:Y* as a link flag(#\_ftn1) (aka "hashtag") attached to the link.

(In Causal Map, the lack of a space behind the colon is essential to ensure this text will be treated as quickfield information.)

Adding quickfield information in this way, even once, ensures that a corresponding field or column ("project") is added to the existing table of links which contains one row for each coding.

In general our cases may be heterogeneous, for example our sources may give us information about different projects but also about different government ministries. So we can need field names like *project* but also *ministry*, etc. Also, although coding caseness is the primary use of quickfields, they can be used for arbitrary tabular data, so in Causal Map we prefer to include in the first part of the quickfield the fact that this particular table is about cases, for example by prefixing with "case-", for example by writing: *case-project:Y* rather than just *project:Y*.

In Causal Map, as we are coding *causal claims*, quickfields can be created for factors as well as for links. It is possible to use the same and/or different conventions for creating case data in factors as in links. So we can store case information in the factors table or the links table. Which is preferable?

As a link contains information about the *entire* coding (not just about one or other of the factors), the default answer to this question is to use the links table, by adding appropriate quickfields to the links flags aka hashtags: *this* (entire) causal claim applies to *that* case. (In the past we have recommended

using factor labels by default but I think this was a mistake.) However there are occasions when it might make sense to use quickfields as part of either or both factor labels too, as appropriate, when the caseness can be conveniently understood as part of the specific meaning of the factor. For example, we might have a hierarchy of factors which includes different kinds of project intervention, and specific types of sub-types of intervention might be associated with specific projects, e.g. Health intervention implemented; deworming / project:X.

## Adding caseness when we already have a 1-1 relationship between sources and cases

Many current Causal Map studies, and all QuIP studies (with the afore-mentioned caveat about focus groups) fall under variant 4ai: multiple sources and cases in a primarily 1-1 relationship. This means that as the (filtered) links table already imports information from the sources table (actually via the statements table), we can trivially use the information about source which is available in each link to ask and answer questions about cases.

Existing metadata about sources...

Sources table		
Source ID	Case-Project	... (many other fields)
V	V	
W	W	
X	X	
Y	Y	

Table 7

... is already imported automatically into the links table:

Links table			
Link ID	Source ID	Case-Project	... (many other fields)
1	V	V	
2	V	V	
3	W	W	
3	W	W	
4	W	W	

5	X	X	
6	X	X	
7	Y	Y	

Table 8

(This way of working is also possible when we have a many-to-one relationship between sources and cases, i.e. variant 4aii: each source deals with only one case but there may be many sources for each case:

Sources table		
Source ID	project	... (many other fields)
a	V	
b	V	
c	W	
d	W	
...	...	

Table 9

... perhaps our ToR says we need to make meta-evaluative judgements not just on the individual projects in their entirety but also perhaps on all the USAID-funded projects. This kind of table can be useful in that case: we can ask and answer questions like “show me all the links just from USAID-funded projects”. But let’s now return to the simpler example with a 1-1 relationship.)

So if we have a dataset with just four sources, project V, W, X and Y, we can ask for example “show only the links from the case *project X*” in literally the same way as we can ask “show only the links from the source *project X*”.

Now, suppose we are serious about coding caseness beyond the 1-1 relationship as in the example in the previous section, i.e. we have added (even just once) *case-project:X* as a link flag, because source W made a causal claim about project X. That means our links table now contains a column *case-project* which is mostly blank but includes the values X or Y.

Links table				
Link ID	Source ID	project	case-project	... (many other fields)

1	V	V		
2	V	V		
3	W	W	X	
3	W	W		
4	W	W		
5	X	X		
6	X	X		
7	Y	Y		

Table 10

In this way we have two separate columns for caseness, one (project) as defined by the source information, and one (case-project) as narrated by the source as part of a causal claim. In this way we can ask and answer questions like “show only the links about case X, whether the information about caseness comes from the source table or from an explicit claim about another project”. Keeping the columns separate like this is probably advisable, so that we can ask and answer questions like “show only the links about case X, and only where caseness comes from the sources table and *not* from an explicit claim about another project”.

If the analyst decides to use a name in the quickfields which is already a name of a field in the links table (e.g. imported from the sources table), then the information will be combined, like this.

Links table			
Link ID	Source ID	project	... (many other fields)
1	V	V	
2	V	V	
3	W	W, X	
3	W	W	
4	W	W	
5	X	X	
6	X	X	
7	Y	Y	

Table 11

(Actually this is not yet implemented in Causal Map, but could be.)

Combining columns in this way is more convenient but it is not recommended because, in the spirit of qualitative research, although of course we want to make useful inferences, we are more interested in describing and exploring and are wary of jumping to conclusions. So we will appreciate for example the ability to show maps in which, where a link expresses a claim about a case other than the source, this information is explicitly printed on the link.

## Qualitative caseness

In real-life projects, especially in variant 4b/ii above, where sources and cases are not 1:1 and where the cases are not clearly defined, at least not in advance, we are most likely to be dealing with *qualitative caseness*: claims involving caseness are made in ways which would frustrate a quantitative researcher, such as

1. “I think *most of the other* projects too were also delayed by strikes”
2. “I think the crops of *some of the other* farmers, especially the smaller farms, were affected by drought”
3. “In *all of the elementary schools*, the new equipment really increased involvement in sports, but *some of the teachers* found it hard to deal with”

Again, using quickfields (or any other similar method) can allow the analyst to quickly implicitly add a new column, or add to an existing column, to define cases on the fly. The contents of the cells in these columns can be any number or piece of text. So we can add, for example, “cases-schools:all-elementary-schools” for the first claim in example 3 above and “cases-teachers-in-elementary-schools:some” for the second. This won’t usher in a new era of qualitative mathematics, but it will help us take some systematic notes to assist our exploration of the data. In Causal Map, it is even possible to put numerical values into a quickfield, for example “cases-teachers-in-elementary-schools:60”; we could interpret this 60 as a percentage and then ask question like “show me the map including only claims applying to 50% or more of the teachers”. However this kind of filter is probably most useful for exploration rather than for making definitive inferences.

## Homogenous and heterogeneous cases

We saw above that with both homogenous and heterogeneous sources, additional information for each source such as gender or location can be provided *a priori* (before coding begins). In variant 4a this can also be done for *cases*, where caseness is clearly defined for each sources simply by providing additional columns on caseness for each source. This way, we can even add additional data for each case a-priori, before we start coding:

Sources table			
Source ID	project	Location of project	... (many other fields)
a	V	Ghana	
b	V	Ghana	
c	W	Nigeria	
d	W	Nigeria	
...	...		

Table 12

But it is less obvious how this would happen in variant 4b, where the cases themselves only arise during coding, or where the cases are not just sources. As we might need to add additional information for cases in variant 4b, there is a good argument for extending the data definition of a Causal Map file to add the possibility of *an additional table for cases*. This might be important for example for coding a meta-evaluation referring to a set of projects, each of which has for example a donor, a location, etc. This additional table would probably not be filled in directly during coding but as a separate step.

Then the analyst could link to these cases using link flags as described above, and it would be possible to ask and answer questions like “show me the map filtered only for projects in Ghana” just as one can already set filters on quickfields which are explicitly mentioned in the link flags.

Cases table			
Case ID (project)	Location	Donor	... (many other fields)
V	Ghana	USAID	
W	Nigeria	USAID	
X	Nigeria	DFAT	
Y	...	...	

Table 13

## Interim conclusion

Most of this is not unique to causal mapping but may be faced by almost anyone doing applied research. The typical working evaluator is quite likely already finding that variant 4b best reflects their daily work. The point is that causal mapping, although specialised for working with causal claims (rather than just

any kind of information) does have tools specifically designed for epistemic aggregation: synthesising what multiple sources believe about multiple things. In the rest of this draft I want to firstly see if causal mapping can address these challenges.

## Part 2: Cases and contexts

### What's a context?

To continue the example of a portfolio meta-evaluation, what do we do with claims like these:

- “The health activities did improve maternal health at the beginning, but by the second phase they didn't seem to be making any difference”
- “The drug was helping the children with their attention problems, but only in a school context, not at home”

The problem with these claims is that if we had perhaps been trying to treat *projects* or *interventions* as cases, now we seem to have claim which only works within part of a project, namely, the early part of it, or part of each clinical case, namely the school context.

Our notion of what constitutes a case starts to fragment: do we have to start treating “parts” of a project, or “aspects” of children's behaviour, as something like as separate sub-cases? Will this fragmentation ever end?

You could say: what really matters is contexts. Cases are just a concrete, realistic approximation to contexts. When we say that 10 cases are covered in this study we can think of e.g. 10 concrete things like projects, teachers, farmers, etc. Whereas a context might be anything: the life-world of two farmers before Covid-19 as contrasted with the life-world of the same farmers after Covid-19: two different contexts.

This question (“what contexts do we have here?”) matters for many reasons – for one, as good qualitative researchers we simply want to correctly describe not only the contents but the structure of the phenomena we are studying. But also, and critically, when doing causal mapping we would like to make transitive conclusions.

### Transitive conclusions and contexts

Causal mapping has tools to help evaluators with many kinds of problem, but probably the most important is tracing the evidence for the influence of one factor or set of factors (X) on another (Y). We have called the map which traces this influence “the traced map from X to Y”; (Bougon et al. 1977) called this an “etiograph”.

What we usually really want to know is answers to questions like this: What influence does X have, in particular does it have an influence on Y? (effects of causes) and What are the influences on Y, in

particular is X an influence on it? (causes of effects). Causal mapping can answer these questions, although prefixed by “*According to these sources ...*”.

Most importantly, we need to avoid what we have elsewhere called the *transitivity trap*. Suppose we learn:

- “When the pig farmers sold a pig, they had increased cash available”
- “When wheat farmers had more cash available, they bought more seed”

Whatever we do, we want to be careful to avoid drawing conclusions like this:

Sell a pig → more cash available → buy more seed

... something which we know is wrong because the first part applies (only) to pig farmers and the second part applies (only) to wheat farmers.

To be sure, if we see that transitivity might break down we could cope with that perhaps by writing the context into the factor labels ...

Sell a pig (pig farmers only) → more cash available (pig farmers only)

... or by using link flags:

Sell a pig (pig farmers only) → *more cash available (pig farmers only)* [= pig farmers only]

(In QuIP reasoning, the rule, though perhaps not always consistently applied, is that we should in any case only ever draw transitive conclusions within individual cases. That means potentially losing data, but it does keep things simple.)

Can we go beyond this ad-hoc solution?

## Drawing transitive conclusions across contexts and sources

The key idea I want to put forward is this: we can be sure transitivity works (only) inside each specific context, because that’s what a context is: in causal mapping, a set of causal claims are from the same **context** if and only if the transitivity rule applies to them.

Every causal link in causal mapping has both a source and a context attached to it. In QuIP individual interviews, the source is just the respondent and the context, where not otherwise specified, is “this respondent, with respect to changes in the last three years”.

When tracing influences, the transitivity rule says:

if source S tell us that B → C in context M, and C → D in context N, then B → D only in the context which is the intersection of B and D. So if B → C applies only to cities in the North, and C → D applies only to border

cities, B à D applies only to border cities in the North. Obviously, if the contexts don't overlap (pig farmers versus wheat farmers) then the conclusion is invalid because it applies nowhere.

We know that S believes the antecedent, have we also concluded that S actually believes, or would at least agree to, B à D? Is this an epistemic conclusion? Are we making conclusions about what someone has to believe? Are we psychologists? *Yes*, because this is a quality standard for the evidence from our sources. QuIP interviewers should interview in such a way that they sometimes actually make such checks (“You told me B à C and C à D, just checking you would agree that, at least indirectly, B à D?”) and even when they don't actually check, they should be reasonably confident that the evidence has been collected in such a way that the source would in fact agree to such a question. Similarly we can assume that the contents of a report or published article are consistent in this way.

What about transitivity across sources? If B à C in context M according to source S, and C à D in the same context M but according to source T, we can conclude B à D in the context which is the intersection of B and D, namely M. But claims in causal mapping are always claims made by someone; there is always a source attached to a claim. So what is the source for this claim B à D? It is a combination of S and T; neither of them need to believe B à D. So rather than saying “B à D according to some assemblage of information from S and T” we had better read simply "*there is evidence that B à D*" – even if no one source actually believes this.

So the complete transitivity rule is:

If

B à C in context M according to source S

and

C à D in context N according to source T

we can conclude:

there is evidence (assembled from different sources) that B à D in the context which is the intersection of B and D.

Causal mapping is not designed to give us rules to conclude causation, i.e. B à D itself (in some context), only that *there is evidence (assembled from different sources) that B à D* in some context.

I know, you're thinking “this seems a bit dodgy to me, assembling claims from different sources, perhaps source S was referring to something slightly different, under different conditions, from source T, so perhaps we aren't warranted to conclude B à D”. But that's what context is for. Your doubt is justified but it's a doubt about context. Indeed we can only conclude B à D in the context which is the intersection of the original contexts. If you have a specific doubt (“I guess S was only thinking about the North part of the

region”) then that needs to be added to the context, making it more restrictive and making the overlap potentially smaller.

## Last-minute principle

And of course we repeat the reminder that this kind of synthesis is a synthesis about evidence for causation, not about causation. What causal mapping does is not to make the leap from cognitions to facts for you, it's to gather together and summarise information about what stakeholders *think* about causation and deliver it back to you as the evaluator so that, if you wish, you can make that leap more easily yourself. Causal mapping tries not to make evaluative conclusions (or even do calculations with evaluative implications) for you, it leaves them until the last minute, for you the evaluator to make.

## What kinds of conclusions can we draw from this kind of data?

Maybe a causal mapping report can contain sentences like these:

“B1 à D, i.e. B1 does causally influence D, according to 12 from 24 different farmers, each reporting on their own context; also there are claims from 3 farmers that B1 à D applies to everyone and there is assembled evidence from 2 more farmers that B1 à D applies to most people; also one of three relevant published reports claims B1 à D but more in the past than now. This is substantially more evidence than for the influence of the alternative intervention B2 on D. Moreover, most of the 12 farmers individually reported a variety of different routes from B1 à D which were robust in the sense of minimum cut; and the overall traced map from B1 à D is much more robust than from B2 à D”.

Causal mapping can help us make and formulate this kind of conclusion, and in particular make different conclusions of this kind comparable.

## Tracing influences in Causal Map

Here I'll just review what we have at the moment at Causal Map in terms of tracing influences.

### Minimum cut and source robustness

We have elsewhere put forward these ideas:

- The "**robustness**" or “minimum cut” measure, which gives an overall, independent score for the “resilience” of a traced map: what is the minimum pieces of evidence/quotes which would have to be removed for there to be no connection. The minimum cut measure gives a good idea of how rich or strong the argument is from X to Y, but it is vulnerable to a form of the transitivity trap, in which individual sources or sets of sources mention separate parts of the map but none (or only a few) mention the entire path. Of course in some studies this is inevitable, for example a forestry study where different stakeholder groups each only know about some specific part of a value chain.
- For this reason we also offer this alternative measure: the “**source thread count**” of the **influence of X on Y** (aka the “traced map from X to Y”) as the number of sources which mention each

part of any path from X to Y (regardless of the causal route taken)(#ftn2). *At the moment in Causal Map, this is now implemented as an add-on to the “tracing paths” function. In contrast to robustness, which is presented as a table, trace\_threads adds fields to the factors and links tables which can be used e.g. in interactive or print maps to show the source robustness from X to \_any Y. Influence can now be traced up or down a map. So we can ask “how many sources mentioned a path to factor Y which began with X” or “how many sources mentioned a path from factor X which ended at Y”. When bundling links, the thread count, the number of threads in each bundle (value=count:threads) can be printed on the links or used to colour or filter them.*

We also note that source robustness is only part of the story:

- it does not distinguish between traced maps in which the sources taken together tend to repeat the same small number of paths repeatedly and traced maps in which each source tends to mention different paths. In the latter case we could say that the paths were more resilient but the evidence for any individual path is thin, and vice versa.
- it does not distinguish between maps in which individual sources tend to mention just a small number of possible paths each, and maps in which sources tend to mention multiple paths, perhaps with multiple evidence for each link. For this reason we also introduce the **extended robustness** measure: the sum of the minimum cut scores for each source. So in traced maps in which each source only mentions one path, and at least one link in the path has only one piece of evidence (one quote), the extended robustness is equal to the robustness, as each source scores only 1 if it has any path, and 0 if it has no path; otherwise, the extended robustness is larger than the robustness. Similarly one could calculate median and mean robustness of all sources, or the median and mean robustness of only those sources with any path, and so on.

We could of course define one single overall metric which sort-of combines these different ideas, but as causal mapping is a qualitative method, we are reluctant to do that, as the right metric and its interpretation depend on the specific task.

We should also note that in the Causal Map functions, both thread count and robustness are both actually more general than described so far: it is also possible to ask about any field in the links table, not just about source\_id, for example how many *villages* or *cases* mentioned a complete link from X to Y.

Another difference between robustness and thread count from a set of factors X to a set of factors Y as currently implemented is that robustness results are available as a table showing the robustness of the influence of *each* factor in set X to *each* factor in set Y, whereas upstream source robustness shows the influence of *all* the factors in X on *each* factor in Y and downstream source robustness shows the influence on all the factors in Y of each factor in X. In other words, in the robustness table you can see the individual robustness of the influence of each factor on each other, whereas in a map showing thread counts you can label outcomes to show the total number of threads from *all* the drivers at once, or vice versa.

## Case robustness?

So the features currently available in Causal Map can deal with sources but only partly address the need to trace transitivity across cases and/or contexts.

Taking into account the first part of this draft essay, we could also define thread count for cases: the **“case thread count” of the influence of X on Y** is the number of cases for which there is a claim (regardless of which sources the claims come from) underlying each part of *any* path from X to Y (regardless of the causal route taken). This can already be calculated and displayed with Causal Map. This generalisation might seem just trivial or confusing for QuIP studies but is essential for studies which do not fall under the 4a/ii variant. However once we get away from the QuIP world, there is no special reason to suppose that counting case threads is free of transitivity problems. It might be, because at least our different sources are talking about the same case; but nevertheless the information comes from different sources.

A corollary of this thinking is, at least in a QuIP study: usually we can be fairly confident that transitivity doesn't break down within the evidence given by a single source: *one source = one context*. If we see that transitivity might break down even within a single case ...

- The drug increased the child's ability to pay attention at school
- At home, her parents praised her whenever she paid more attention

... where we do not want to conclude:

Drug → child pays more attention → parents praise child

... we are obliged to ensure using some other means that transitivity remains broken, perhaps by writing the context into the factor labels.

Drug → child pays more attention (at school only)

... or by using link flags:

Drug → *child pays more attention* [= at school only]

If we use factor labels, there is no danger of falling into the transitivity trap, because the factor child pays more attention (at school only) is literally a different factor from child pays more attention and therefore we cannot take this path from Drug to parents praise child. Using link flags (or generally attaching any kind of additional information or warning to a link) only guarantees not falling into the transitivity trap if we include some extra rules like “be careful following links with this kind of flag”.

## Avoiding the transitivity trap when zooming out

One more issue is that there are special and more pernicious forms of the transitivity trap when zooming out from a hierarchical map. It is really not obvious that if we zoom out from:

Hand-washing intervention → improved health behaviour; hand washing

And

Improved health behaviour; not smoking → less lung disease

We should be able to conclude

Hand-washing intervention → less lung disease.

(Arguably, conventional statistics which use sets of indicators for latent variables is also subject to this trap.) What's really confusing in Causal Map is that our robustness tables do not do anything about this trap, whereas thread count does take special measures to avoid it: it actually remembers the paths as they were before zooming out and traces influence across the old paths.

## Conclusion and takeaways

My aim here is really to contribute to our discussion about sources, cases and context and what to do about them when coding and analysing. The essay peters out here because I want first to experiment with actually coding a file using these kinds of ideas more explicitly than I've done in the past, before I make stronger suggestions on what exactly we should do.

So I think when designing and training the procedures for eliciting causal claims, e.g. interview protocols and training, we should:

a) Try to ensure that transitivity applies to all the claims from this source, i.e. if you were to assemble their causal map in front of their eyes, they would assent to all the causal paths and agree that if they said B → C and C → D then they also assent to B → D. If not, add contextualising information. (I don't suggest this should be done formally, it's just something to bear in mind.)

Also I think when designing the coding procedures for a file, we should:

b) be aware of whether we have homogenous or heterogeneous sources. If the latter, combine the information into a sources table like Table 5 above, and consider reporting the different types of sources separately.

c) Think hard about the relationship between sources and cases as in Table 6 above. Are they 1-1?

d) If the sources and cases are separate but related (e.g. several sources for one case), construct the sources table carefully as in table 9, so that information about cases as well as about sources is available to filter and format maps.

e) Even cases and sources are 1-1, do the sources nevertheless sometimes start talking about other cases which we need to actually code (rather than relying on metadata) and if so how are we going to code

this (“coding caseness”)? Are we going to be coding only using cases which already appear in the sources table (as in QuIP studies) or do we have completely new sets of cases?

- f) If we are coding caseness,
  - a. do we need to add metadata about these cases? (this is in fact not yet possible in Causal Map, see Table 13).
  - b. consider using quickfields like case-project:UNICEF (leaving open the possibility of heterogeneous cases, with formulations like case-ministry:health as well).
  - c. Consider using qualitative coding like case-farmers:some.
- g) Be aware that whether we think about cases or sources, additional conditions may come into play which put caveats on the causal claims. Fundamentally, we can decide whether to treat cases (or even sources) as complete contexts, i.e. for each case (or source), transitivity works for all the causal claims. Where this breaks down (e.g. a source says something which does apply to themselves but only under certain conditions, or which only applies now but not recently) we must specify this by using additional factor or link flags.

When analysing files we should:

- h) If explicitly thinking in terms of cases instead of or as well as sources, consider analyses which explicitly address cases, e.g. filter maps to only show certain cases.
- i) Where we have specified special conditions/context for certain claims which do not apply to the whole source or case,
  - a. if we used factor flags to do this we can't fall into the transitivity trap, so nothing to worry about;
  - b. if we used link flags, we have to remember that transitivity only works in the intersection of the contexts.
- i. By default, don't include links with context flags in ordinary analyses
- ii. Conduct special analyses with contexts like “context-time:future”.
- j) When drawing transitive conclusions, i.e. when looking at traced maps in order to answer questions about effects of causes and causes of effects:
  - a. When conclusions are assembled from multiple sources, consider reporting this with formulations like “*there is evidence (assembled from different sources) that B à D in some context*”.
  - b. Use the updated Trace Paths button in Causal Map (partly completed as of 11/8/2022) for all your path tracing, optionally specifying if you want to also trace threads and/or calculate robustness.

- c. Consider using robustness to describe and compare how much traced maps rely heavily on just a small number of pieces of evidence (low robustness).
- d. Consider using source robustness to display robustness on a per-source basis
- e. Consider using source thread counts (or case thread counts) to avoid falling into the transitivity trap, and consider excluding all links with contextual caveats from main analyses and analysing these separately.
- f. Think carefully about tracing threads and calculating robustness on maps which have been zoomed out.

Also, thread tracing should probably only count threads from the same context.

Marking links as only applying in a particular context (and then only counting threads if they all share it) is a better realisation of interaction effects.

## References

Barbrook-Johnson, Pete, and Alexandra Penn. 2021. “Participatory Systems Mapping for Complex Energy Policy Evaluation.” *Evaluation* 27 (1): 57–79.

Barbrook-Johnson, Pete, and Alexandra S. Penn. 2022. “Comparing, Choosing, and Combining Systems Mapping Methods.” In *Systems Mapping: How to Build and Use Causal Models of Systems*, edited by Pete Barbrook-Johnson and Alexandra S. Penn, 161–77. Cham: Springer International Publishing.

Bougon, Michel, Karl Weick, and Din Binkhorst. 1977. “Cognition in Organizations: An Analysis of the Utrecht Jazz Orchestra.” *Administrative Science Quarterly*, 606–39.

Wilkinson, Helen, Dione Hills, Alexandra Penn, and Pete Barbrook-Johnson. 2021. “Building a System-Based Theory of Change Using Participatory Systems Mapping.” *Evaluation* 27 (1): 80–101.

---

(#\_ftnref1) We’ve used the word “hashtag” extensively for link flags but it’s a really confusing name.

(#\_ftnref2) There could be an alternative measure of “source robustness” of the influence of X on Y: the minimum number of sources whose evidence would have to be removed in order for there to be no evidence for any path from X to Y. While the two are probably related, they are not the same, because in the map  $X \rightarrow V \rightarrow Y$ , where the two links come from different sources, the actual source robustness is 0 whereas this alternative measure would be 1.

## Related

- [chapter intro](#)

---

## References

- Barbrook-Johnson, & Penn (2021). *Participatory Systems Mapping for Complex Energy Policy Evaluation*. <http://dx.doi.org/10.1177/1356389020976153>.
- Barbrook-Johnson, & Penn (2022). *Systems Mapping: How to Build and Use Causal Models of Systems*. Springer International Publishing. <https://link.springer.com/10.1007/978-3-031-01919-7>.
- Bougon, Weick, & Binkhorst (1977). *Cognition in Organizations: An Analysis of the Utrecht Jazz Orchestra*. JSTOR.
- Wilkinson, Hills, Penn, & Barbrook-Johnson (2021). *Building a System-Based Theory of Change Using Participatory Systems Mapping*. SAGE Publications Ltd. <https://doi.org/10.1177/1356389020980493>.