



## Causal Loop Diagrams

**Abstract** This chapter introduces Causal Loop Diagrams. We explore what exactly Causal Loop Diagrams are, describe how you can use them, take a step back to consider common issues and ‘tricks of the trade’, as well as present a brief history of the development of the method. This chapter can be viewed as a companion to Chap. 8 on System Dynamics; these two methods are closely related. Causal Loop Diagrams emerged from Systems Dynamics practice, and though it is a systems mapping method in its own right now, it is still often used as a stepping-stone to the development of System Dynamics models.

**Keyword** Causal Loop Diagram

Let’s get the elephant in the room out of the way. The obesity system map (Government Office for Science, 2007), which you have likely seen, was produced for the UK government back in the mid-2000s and is a Causal Loop Diagram (CLD). It was, and still is, one of the most high-profile pieces of systems mapping work ever done and has been very successful and influential by most measures. However, it is also often derided or held up as an example of a ‘horrendogram’ (Penn, 2018); indicative of a form of visualisation which renders problems intractable and overwhelming.

Even without the obesity example, CLDs are a no-brainer to include in this book. They are a well-used approach to systems mapping, with roots

going back to the 1970s at least, and a strong connection to System Dynamics. The method sits roughly in the middle of the qualitative-quantitative spectrum of systems mapping methods, perhaps more towards the qualitative side, but hints at dynamics in a system. It focuses on feedbacks as a key component and organising structure for complex systems, but also brings the ability to include all sorts of concepts and variables. This means it can facilitate the production of focused maps, which strongly hint at the dynamics of systems, but also much larger inclusive ones, which allow us to see the big picture and the interconnected nature of systems. CLDs also have one of the most consistent and appealing visual styles of any of the qualitative methods in this book. Arguably, this started with the influential ‘Foresight’ work (done in the UK), which showed what proper resourcing for design could do for developing high-quality visualisations. This tradition has been extended with contemporary software tools, such as Kumu, which place a high value on slick and appealing visualisations. Finally, it is worth noting, CLDs are a natural stepping-stone to simulation methods such as System Dynamics because of their focus on feedbacks and strict use of variables for nodes in the map.

The rest of this chapter will follow the same structure as the other methods’ chapters in this book. First, we describe exactly what a CLD is. Next, we describe how the method is ‘done’. Then we consider common issues and some ‘tricks of the trade’ for using the method. Next, we reflect on what the method is good and bad at, before exploring a brief history of its development. We finish with some useful resources and tips for getting started with the method yourself.

## WHAT IS A CAUSAL LOOP DIAGRAM?

CLDs represent a system in three basic elements: boxes, connections, and feedback loops. The boxes, or nodes, represent variables in the system; these can be anything as long it makes sense to think of them going up or down over some scale. The connections, or edges, represent causal influence, from one node to the other; either positive (i.e. they increase or decrease together) or negative (i.e. they change in opposite directions, if one goes up, the other goes down, and vice versa). So far, fairly similar to Fuzzy Cognitive Mapping or Participatory Systems Mapping. The third element is what makes CLDs more unique. The maps always show and focus on feedback loops, both in the construction of the map and in its visualisation. Loops are made conspicuous by the use of curved arrows to

create circles. The loops are also sometimes colour coded to highlight them or are annotated with small arrows and ‘+’ or ‘-’ symbols to highlight if they are reinforcing (positive) or balancing (negative) feedback loops.

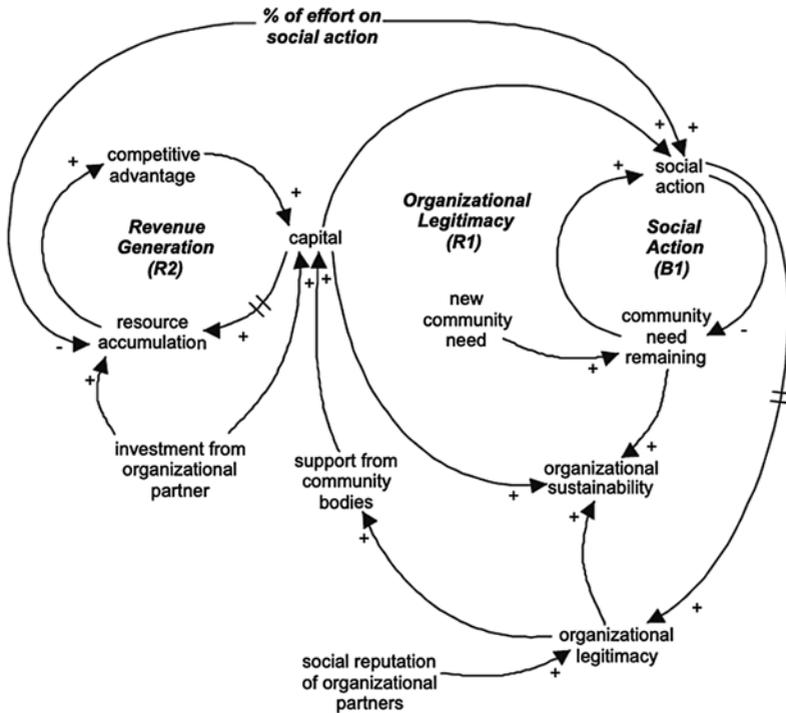
The feedback loops are usually focused around a ‘core system engine’, which is a set of nodes that are the core of the system. These are often visualised more prominently than other nodes in the map. The feedbacks strongly hint at dynamics in the system. It is common to use the maps to think at a slightly higher level than individual nodes and edges, bringing together the handful of feedbacks in the system to think about how these might play out together. The focus on feedbacks means the CLDs are a relatively disciplined way of looking at a system, which places the existence and effects of feedbacks at its core. It is common to see both simple (perhaps five or ten nodes) and more complex CLDs (dozens and dozens of nodes).

Let’s look at some examples. Figure 4.1 shows an example of a simple CLD. We can see how nodes and edges are connected with positive and negative connections, and how the feedback loops of ‘revenue generation’, ‘organizational legitimacy’, and ‘social action’ are emphasised through their positioning, and additional annotation. The ‘R1’, ‘R2’, and ‘B1’ refer to reinforcing and balancing loops.

The next example, shown in Fig. 4.2, is the well-known and much larger CLD of the obesity system in the UK. Here, nodes are in boxes and colour coded by theme, and the connections are shown with solid lines (positive connection) and dotted lines (negative connection). The core engine of ‘energy balance’ and its feedbacks are highlighted in the centre.

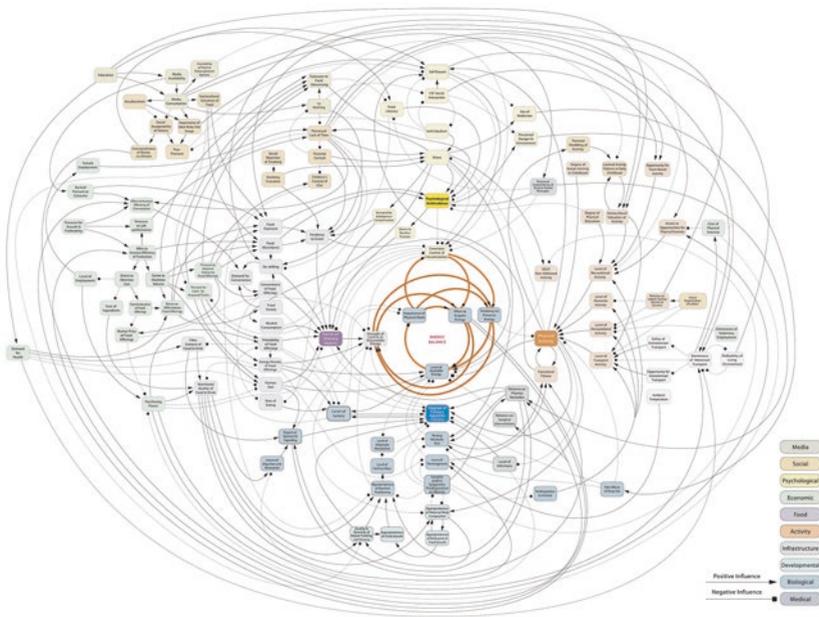
CLDs are used in a variety of ways. Most fundamentally, as with many methods in this book, they are a way of surfacing, visualising, and exploring mental models. Exploration of the full map is done in many ways, but qualitative analysis of the core engine and feedback loops is a common focus. CLDs are often a precursor to System Dynamics models (Chapter 8); used to begin the modelling process in an intuitive way, before the conversion to stock and flow diagrams and differential equations.

The maps can be built from all types of information, and in participatory modes, and the mix between these is fairly even in the literature. There is a lot of variety in the ways they are built and the exact purpose they are put to. However, there is less variety in what is actually constructed; the use of clear variables, arrows for influence, and focus on feedbacks is consistent. There is a lot of quite prescriptive guidance available,



**Fig. 4.1** A Causal Loop Diagram of the tensions between business activity and social action in a social enterprise. Note, double bar (//) symbols indicate a time delay. Three feedback loops are emphasised (B1, R1, and R2). Source: Moizer and Tracey (2010)

which helps to maintain this consistency. However, it is worth noting, there does appear to be a divide between those who use CLDs as a stand-alone method and those that use them as a stepping-stone to System Dynamics. The practice of each group does tend to have subtle differences, with the former being more inclusive and flexible in their use of the method, and the latter following clearer rules about how to develop a map. There is a little variety in the terminology used to describe CLDs; they are often referred to as ‘influence diagrams’, or simply ‘system maps’. There is also a small group of papers which use the term ‘Participatory



**Fig. 4.2** A Causal Loop Diagram of the obesity system in the UK. Source: Produced by ShiftN for Government Office for Science (2007). The core feedback loops are emphasised in the centre

Systems Mapping’ to refer to a participatory CLD approach (e.g. Lopes & Videira, 2016).

### HOW DO YOU CREATE CAUSAL LOOP DIAGRAMS?

The process of building and using a CLD is typically iterative. They can be developed in a participatory workshop mode or based on data (typically qualitative) and evidence. However it is done, the process tends to include the following steps:

1. **Collect data and/or evidence:** first you need to know what you are going to build your map from. Maps can be built in a participatory mode, in which case, discussions during workshops will be your data. However, because of the strong discipline needed to focus on

a system engine and feedbacks, decisions are often made by the modeller/researcher. When built from data or evidence, these need to be collected or compiled before mapping can start.

2. **Screen data for variables and connections, create catalogue:** once you have your underpinning data or evidence, you need to screen and extract potential variables and connections from them, and build a catalogue of them. In a workshop setting, this could be a brainstorming exercise, or could be done at the same time as actually drawing the map. The variables should be expressed clearly as things that can go up or down and should be precise. Try to avoid developing ‘container concepts’ such as ‘Technology’ which can mean many different things. If a connection feels as though it is unclear or complex, consider if adding a second will make it easier. For example, we might think a ‘price of good’ variable has an unclear influence on a ‘revenue’ variable (i.e. higher prices equal more money per sale, but potentially fewer sales), but if we add another variable, ‘sales’, we can incorporate this effect separately and remove the uncertainty in individual links.
3. **Hypothesise the structure and content of the core system engine:** this is the first CLD-unique stage. Assuming you already have a strong sense of the purpose of the mapping project, and you understand users’ interests, this is the stage at which you start to focus in on what might be the ‘core system engine’. There is no simple way to decide this. It may be that there is one defining variable at the centre, and that is enough; or you may have a handful of key variables which are of most importance to people in the system, or which interact in ways which drive system behaviour. Alternatively, you may have an idea or narrative about what sort of dynamical behaviour is at the heart of the system. The engine will be the focus of the map as a whole and will be one of the main focuses for exploring feedbacks. At this stage, you will need to identify which variables are included in the engine and begin to develop ideas about how they might be connected. This might be as far as you hope to get in a first workshop, before stopping and starting the next stage in a desk-based research mode. This stage can be difficult, it relies on craft and judgement more than objective fact or technical skill. You may find it useful to refer to the ‘system archetypes’ developed by Kim (2000), which outline common and intuitive types of core system engines; they are: ‘drifting goals’, ‘escalation’, ‘fixes that fail’,

‘growth and under-investment’, ‘limits to success’, ‘shifting the burden’, ‘success to the successful’, and ‘tragedy of the commons’.

4. **Develop first diagram with core system engine and twenty nodes:** now you are ready to build the first diagram of the core system engine and the variables immediately around it. This will be a stage with many iterations and repeats as your ideas crystallise. Once you have the engine down, you can start to build in a similar way to other methods, either brainstorming (or extracting from your catalogue) the top twenty variables that are influenced by, or influence, the core, and connecting them; or by more organically building one variable at a time. There are some excellent and detailed guidance for this stage at [www.thesystemsthinker.com](http://www.thesystemsthinker.com), including how to deal with ambiguities, including delays in influence, and naming feedback loops. Twenty nodes seem to be a useful limit here, beyond this the core becomes too unwieldy.
5. **Verify and amend:** this stage is arguably the most important, but the most underestimated. Once you have a first version of the core system engine and surrounding nodes, you need to expose it to as much critique, feedback, and commentary as possible. This may be with stakeholders who you are working with, with clients or users, or with other researchers and modellers. You should amend the core system engine considering this feedback and your aims or create different versions if necessary.
6. **Expand model as needed:** once the core feels more settled, you can expand the map if needed. It may be that you want a simple map, and the core you already have is enough, in which case you can skip this stage. Or, you may want to build a much larger map, in which case you should expand and then iterate through another ‘verify and amend’ stage.
7. **Analysis and use:** once you have the map, there are several ways it can be used and analysed. Purely by its existence it makes clear your, or your stakeholders, mental models. You might want to do a thematic analysis of it, either in written form or in a workshop setting with stakeholders. You may want to explore it qualitatively, that is, investigate how the various feedbacks in the map might play out and ask questions such as ‘is the system in a stable lock-in or might there be reinforcing feedbacks which cause runaway change?’. It is common to create mock-up plots of variables through time to explore these. Finally, if you are using the map as a starting phase for a

System Dynamics project, you are now ready to start the process of converting it; there is no simple recipe to follow for this. At this stage, it is also common to finalise the visualisation of the map; many CLD projects put a lot of value in a well-designed and appealing final map. You may want to consider asking a designer to help you with this.

### COMMONS ISSUES AND ‘TRICKS OF THE TRADE’

Most of the issues that are more specific to CLDs tend to revolve around the use of feedbacks. Constructing feedback loops requires a different way of thinking, and perhaps a little more care than other parts of the map, or other types of maps where we tend to work solely in variables and connections. It brings an additional layer of easy mistakes to make. Fundamentally, this is because we are building a group of important and interacting variables and links together. This means we need to hold more information in our minds at once, and imprecisions can creep in more easily. Guidance on CLDs often provide tips such as labelling loops or making the goals of balancing loops clear (i.e. including a variable which expresses a goal and linking it to a loop). These are specifically aimed at helping loop construction go well.

When it comes to interpreting and using CLDs, it is easy for the feedback loops to be misunderstood or skipped over. Many people will find it hard to link together multiple loops and think about how these may interact through time. It can be useful to create mock-up plots of variables through time to explore how the system might behave. Others, rather than finding loops difficult, may simply not realise their significance in a CLD and ignore them. This is one of the easiest ways to miss one of the key values of CLDs. Even when loops are noticed, understood, and discussed, it can be difficult to turn this dynamical view of a system into something usable or ‘actionable’. The map will give us the language and understanding on an individual level, but problem-solving cultures in organisations can get in the way if they revolve around instrumental and ‘button-pushing’ mindsets. This is an issue with most systems mapping methods, but it really comes to the fore with CLDs because of the combination of accessibility and exploration of dynamics. The system archetypes described by Kim, for example, are accompanied by suggested ways to reframe problems or redesign situations to help change dynamics, but this certainly requires decision-makers taking a systems view. One more

straightforward action might be identifying simple positive/reinforcing, or negative/balancing, feedbacks, which might be encouraging or suppressing change that we do or don't want. Disrupting a 'vicious cycle' or encouraging a 'virtuous' one is probably the most intuitive way in which we can think about interaction with system dynamics.

There are a few 'tricks of the trade' to help us deal with these and other issues, these include:

- **Tune the complexity of a map to the visual literacy of your audience:** while we often talk about bringing users, clients, and stakeholders into our thinking on purpose and design of a systems mapping project, we don't normally talk about their visual literacy. CLDs have a good tradition of taking real care over visualisation, and part of this includes thinking about the visual literacy of your audience. If you think your audience will be put off by large maps, or be confused by multiple interacting feedback loops, think about how you can tune your map to this, start small with the engine, and build up. The point should not be to just produce a very simple map, though you may end up doing this, but to also think carefully about how you might frame and introduce it in ways which allow you to keep as much complexity as possible (e.g. introduce it in steps, or colour code and position nodes by themes).
- **Allocate half of your resources to design and communication:** it is natural to focus our energy on creating a map. However, when we are using maps with subtleties beyond face value and that need to be created with a lot of control by researchers (i.e. like CLDs), we really need to invest in our final visualisations and communication of maps to users and stakeholders. Otherwise, we will have created a brilliant CLD which no-one is using or looking at. This takes time and effort, we would suggest thinking about using half of your time or resources on this; it can include a range of activities, from smart visuals, web-based interactive maps, to events to disseminate and use the map.

### WHAT ARE CAUSAL LOOP DIAGRAMS GOOD AND BAD AT?

CLDs' real strength comes from producing a system map that is often visually appealing, that hints strongly at a system's dynamics, but that is still flexible, inclusive, and relatively easy to use. CLDs can make use of all

sorts of information, can be big or small, and offer a clear stepping-stone to producing a more formal System Dynamics models.

On the downside, CLDs can be restrictive with their focus on feedbacks and do tend to put a lot of power in the hands of the researcher, especially around the creation of feedback loops. If feedback loops are not present, or important, in your system, then the method will be less useful. Being in the middle of the quant-qual spectrum means CLDs will also not give you any quantitative analysis. Relatedly, without any quantitative analysis, it can be difficult to meaningfully understand how multiple feedback loops will interact.

### A BRIEF HISTORY OF CAUSAL LOOP DIAGRAMS

CLDs' early history is intertwined with that of System Dynamics. In the early work of Jay Forrester introducing System Dynamics, there is no mention or use of CLDs, though he does use flow diagrams as a mid-point between a verbal description and a formal model of a system. CLDs appear to have emerged later as a way of communicating the design of a System Dynamics model. Goodman's (1975) 'Study notes on System Dynamics' is one of the earliest texts to formally outline CLDs and does suggest they can be used as a tool for model conceptualisation (i.e. in the design stage) too. Simple CLDs were also used in the famous Limits to Growth report (Meadows et al., 1972), which is likely to have raised awareness of them.

In the earliest works, CLDs were primarily used to communicate the design of a System Dynamics model. However, their use has spread beyond this to such an extent that we now think of them as a method in their own right. Moreover, when used in conjunction with System Dynamics, it is more common to use them as a stepping-stone towards a model, rather than as a tool to help communicate them after they have been built.

The 'rules of the game' for how to construct CLDs, and what is included in them, have remained relatively stable through the years. There have been debates around relatively minor issues, such as whether links should be labelled with '+' and '-' symbols or the letters 's' and 'o' (to represent 'same direction' and 'opposite direction'). What has clearly changed is the prominence they are given as a method in their own right, the ways we might build them (i.e. not just from a modeller's design, but from data, evidence, and participatory processes), and the ways we might use and analyse them (i.e. with exploration of the perceived behaviour and interaction of loops).

## GETTING STARTED WITH CAUSAL LOOP DIAGRAMS

CLDs are one of the methods which is relatively easy to get started with. If you just want to have a play and get a feel for the method, we recommend just diving-in and building one, either on pen and paper or with specialist software. There are numerous guides available which will be useful to have by your side. The process of implementing the tips on your map will quickly allow you to understand how the focus of a CLD is different to other types of system maps.

More generally, there are range of resources we would recommend for helping you go deeper into CLDs, including:

- [www.thesystemsthinker.com](http://www.thesystemsthinker.com): this is a website version of the original ‘The Systems Thinker’ publication which has run since the early 1980s. There are hundreds of short and accessible articles on all sorts of topics, including dozens for CLDs. We recommend using the search function to find CLD articles, as there is not a stand-alone category for browsing.
- **System archetypes (Kim, 2000)**: this suite of guides is invaluable for getting into the most important details of CLD. There is a wealth of information in each one, but the eight archetypes, and the discussion around them really gives a sense of what the core engine of a system might look like.
- **Group model building textbook (Vennix, 1996)**: though this book is focused on System Dynamics as a whole, it does consider ‘qualitative System Dynamics’ which includes CLDs, specifically, and is a foundational text for the participatory development of these types of models.
- **Obesity and land use futures examples (Government Office for Science, 2007, 2010)**: finally, it can be useful to see CLDs ‘in action’. For this we would highly recommend reading two examples of their use in the Foresight work we mentioned above; one on the famous obesity map, and one on land use futures. You can avoid the long-form reports if you want and focus on the separate publications specifically about the CLDs work that was done.

You can build a CLD with pen and paper, or on using general-purpose diagramming software, but it will be quicker, and possibly more visually appealing, to use some purpose-built software, such as:

- **Kumu:** this increasingly popular software provides an excellent web-based interface for building and analysing system maps. Though it is flexible, it broadly uses the underlying logic of CLDs and so is well-suited to them. You can build ‘public’ maps for free, but there are fees to use the software to create private maps.
- **System Dynamics software (see Chap. 8):** most System Dynamics software has the capability to create CLDs, for example, iThink/Stella, Powersim Studio, and Vensim.

Hopefully, you now have a strong sense of what CLDs are and how you can use them. Take advantage of the fact that this is one of the easier methods to get started with; have a play around with some software today, or read one of the guides and dive in. Be warned though, you can get hooked easily, and before you know it you will be looking at everything through the lens of feedbacks and core system engines! As a method that is easy to start with, it is possible to go quite far without reading some of the more advanced guides and textbooks; we would advise against this if possible. It can really improve the quality of your mapping work if you check in with these resources regularly; much of their advice will only really sink in once you are up and running. Good luck!

## REFERENCES

- Goodman, M. R. (1975). *Study notes on system dynamics*. Pegasus Communications.
- Government Office for Science. (2007). *Reducing obesity: Obesity system map*. <https://www.gov.uk/government/publications/reducing-obesity-obesity-system-map>
- Government Office for Science. (2010). *Land use futures: Making the most of land in the 21st century*. <https://www.gov.uk/government/publications/land-use-futures-making-the-most-of-land-in-the-21st-century>
- Kim, D. (2000). *Systems Archetypes I*. Toolbox reprint series. Pegasus Communications. <https://thesystemsthinker.com/systems-archetypes-i-diagnosing-systemic-issues-and-designing-interventions/>
- Lopes, R., & Videira, N. (2016). Conceptualizing stakeholders’ perceptions of ecosystem services: A participatory systems mapping approach. *Environmental and Climate Technologies*, 16(1), 36–53. <https://doi.org/10.1515/ruct-2015-0011>
- Meadows, D. H., Meadows, D. L., Randers, J., & Behrens, W. W., III. (1972). *Limits to growth. A report for the Club of Rome’s project on the predicament of mankind*. Universe Books.

- Moizer, J., & Tracey, P. (2010). Strategy making in social enterprise: The role of resource allocation and its effects on organizational sustainability. *Systems Research and Behavioral Science*, 27, 252–266. <https://doi.org/10.1002/sres.1006>
- Penn, A. S. (2018). Moving from overwhelming to actionable complexity in population health policy: Can a life help? *Artificial Life*, 24(3), 218–219. [https://doi.org/10.1162/artl\\_e\\_00265](https://doi.org/10.1162/artl_e_00265)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

