CASCADE INSTITUTE

Interventions for Rapid Global Change

Causal Loop Diagrams: A Short Handbook

Michael Lawrence,
Polycrisis Research Fellow

www.cascadeinstitute.org

Summary

“Systems mapping” encompasses various graphical techniques that help us to think through our mental model of a system, increase our understanding of a system, and communicate our knowledge to others.[1] A causal loop diagram is just one type of systems map, but an especially useful one, because it captures the feedback loops that constitute the basic structure of the system and shape its key behaviours. Causal loop diagrams (CLDs) do not require any specialized knowledge but can be counter-intuitive in some respects. This handbook explains how to read and draw causal loop diagrams, clarifies some common sources of confusion, and offers advice for using this tool most effectively.

Like many other forms of systems diagrams (and network diagrams), CLDs are composed of elements and connections. But unlike many others, CLDs also include feedback loops that connect elements in a circular pattern. This handbook explains each of these three features, provides step-by-step instructions for drawing CLDs, then presents three examples of CLDs that elucidate crucial real-world phenomena.

---

1 For an excellent survey of different system mapping techniques, see: Barbrook-Johnson and Penn, 2022, in the Recommended Resources section below.

Contents

ELEMENTS

Elements are the key variables in a system—that is, the causal factors that are most central to its behaviours. They might encompass "stocks" (such as available resources), "levels" (such as the amount of trust in government), "shocks" to the system (such as extreme weather events), or myriad other relevant causal factors. The one major constraint the CLD methodology imposes on the identification of elements is that these variables must be able to increase and/or decrease. A given actor—say, the Canadian government, for example—would not fit as an element in a CLD, but the Canadian government's power and its budget would both fit, as they can rise and fall over time in relation to other causal factors. Elements are presented as words in a CLD, often inside shapes, and connected by arrows. If helpful, you can use different shapes and colors to distinguish different types of elements.

CONNECTIONS

Causal loop diagrams involve two types of connection: positive causal relationships and negative causal relationships. The Figures below (including those in the sections on positive and negative feedback loops) show elements expanding or contracting in relation to one another in order to illustrate these two types of causal relationships. Actual CLDs, however, do not depict variables as expanding and contracting in this way; they leave it to the viewer to mentally interpret the ways in which variables increase and decrease in relation to each other.

Positive causal relationships are depicted as an arrow with a plus (+) sign above it (Figure 1a) that connects two elements. In positive relationships, cause and effect change in the same direction: when the cause increases in strength, the effect increases in strength (Figure 1b); when the cause decreases, the effect decreases (Figure 1c). The cause always precedes the effect in time.

Figure 1: Positive causal relationships

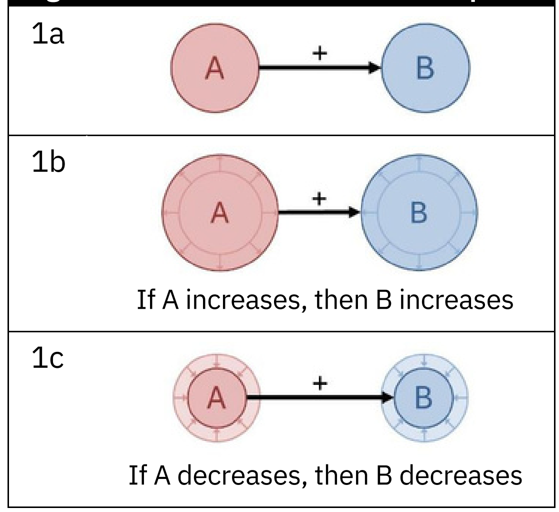

Negative causal relationships are depicted as an arrow with a minus (-) sign above it (Figure 2a) that connects two elements. In negative relationships, cause and effect change in opposite directions: when the cause increases in strength, the effect decreases in strength (Figure 2b); when the cause decreases, the effect increases (Figure 2c).

Figure 2: Negative causal relationships

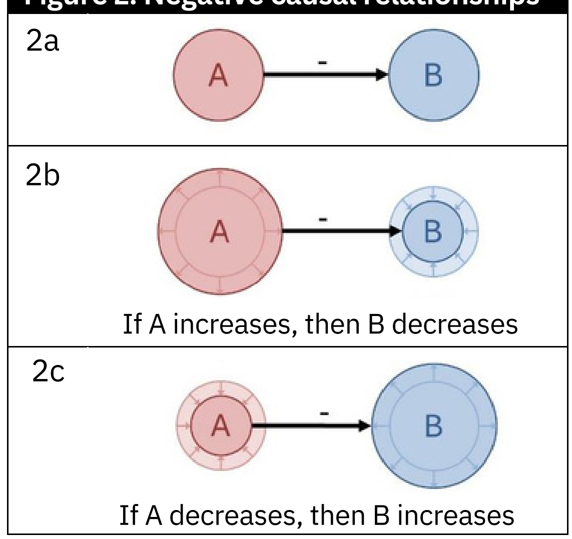

POSITIVE FEEDBACK LOOPS

Feedback loops are chains of positive and negative causal relationships in which an initial cause produces an effect or (or series of effects) that has an impact on the initial cause, thereby forming a causal loop. There are two types of feedback loops: positive feedback loops and negative feedback loops.

In a **positive, or self-reinforcing, feedback loop** (Figure 3a), a change in the cause creates a similar change in the effect, which then intensifies the change in the cause, which then intensifies the change in the effect, and so on. An increase in A (Figure 3b) causes an increase in B (Figure 3c), and that increase in B causes a further increase in A (Figure 3d), which causes a further increase in B (Figure 3e), which causes an additional increase in A (Figure 3f), and so on.

A positive feedback loop operating in an increasing direction produces runaway growth.

Compound interest is a good example of a simple positive feedback loop, where A represents your account balance and B represents the amount of interest you earn annually on the money in your account. Increasing your account balance (Figure 3b) increases the amount of interest you earn (Figure 3c), which increases your account balance (Figure 3d), which increases the interest earned (Figure 3e), which increases your balance (Figure 3f), and so on.

Figure 3: A positive feedback loop working in an increasing direction

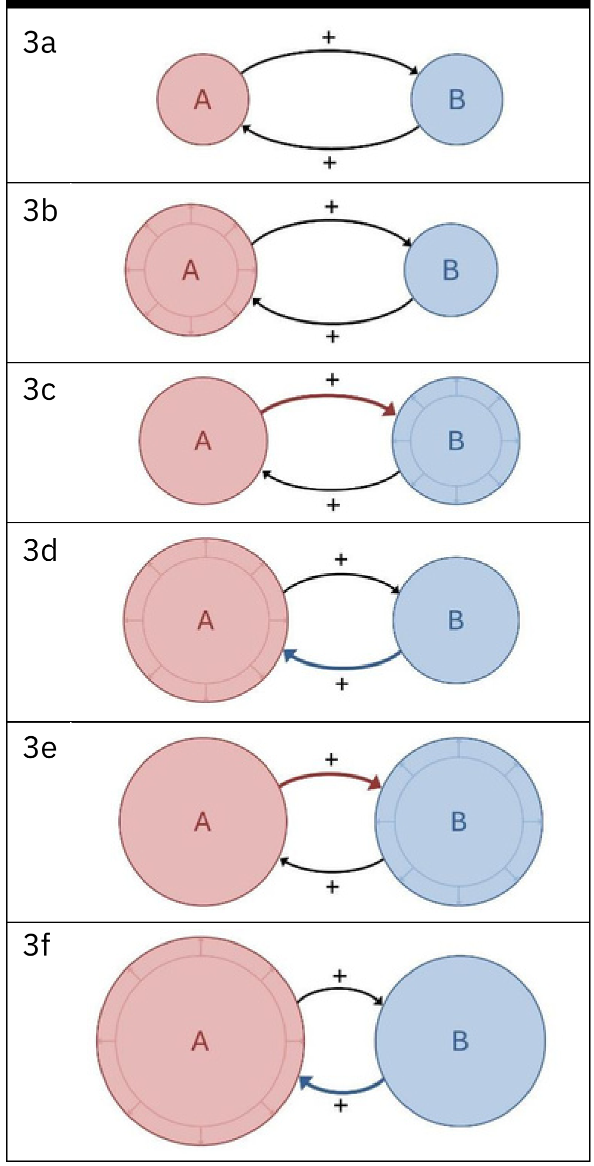

Figure 4 shows the runaway, exponential growth on an initial account balance of $100 with an annual interest rate of 5%. Left alone, the balance grows to over $162 in ten years, about $340 in twenty-five years, and nearly $1150 in fifty years. Even though the interest rate remains constant, the account balance grows in a non-linear manner.

Figure 4: Compound interest

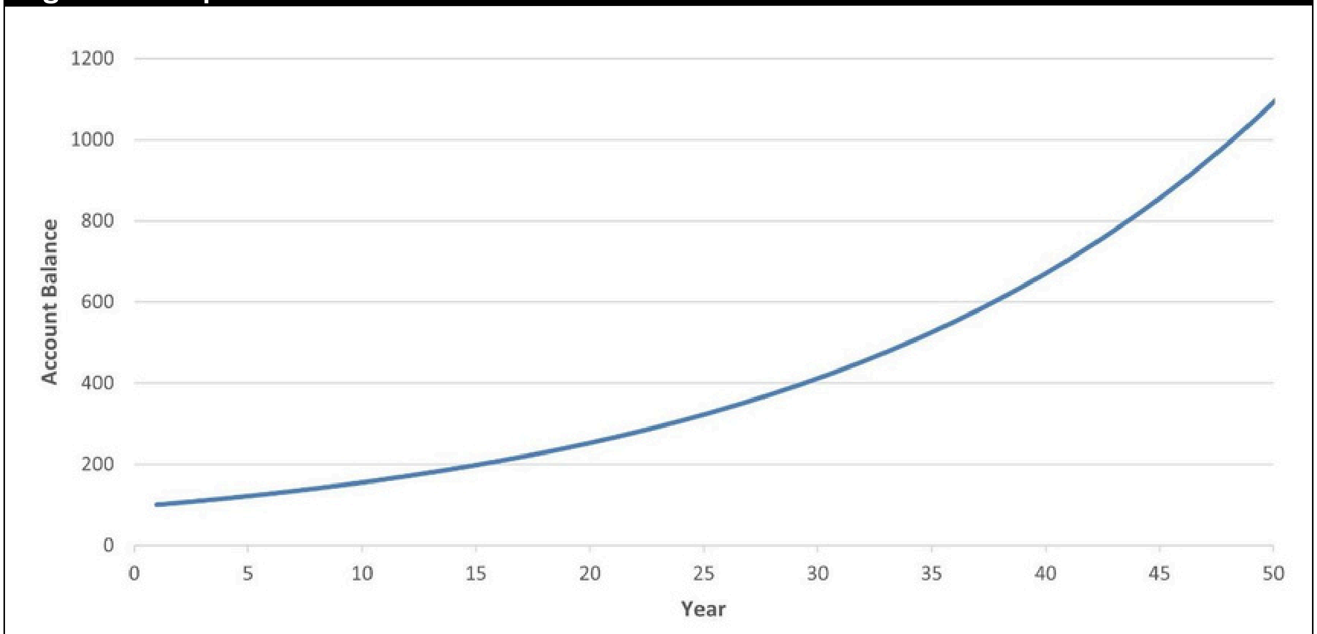

Figure 3 shows a **positive feedback loop** proceeding in an increasing direction; but positive feedback loops can also proceed in a **decreasing direction**, as depicted in Figure 5.

A decrease in A (Figure 5b) causes a decrease in B (Figure 5c), and that decrease in B causes a further decrease in A (Figure 5D), which causes an additional decrease in B (Figure 5e), and so on.

A positive feedback loop operating in a decreasing direction produces a spiral down to nothing.

The loss of trust between two people may unfold in a positive feedback loop working in a decreasing direction. Imagine A represents the trust I show in you, and B represents the trust that you show in me. If I show you less trust than before (Figure 5b), it may cause you to show less trust in me (Figure 5c), which causes me to show less trust in you (Figure 5d), which causes you to show less trust in me (Figure 5e) and so on until we do not trust each other at all.

Figure 5: A positive feedback loop working in a decreasing direction

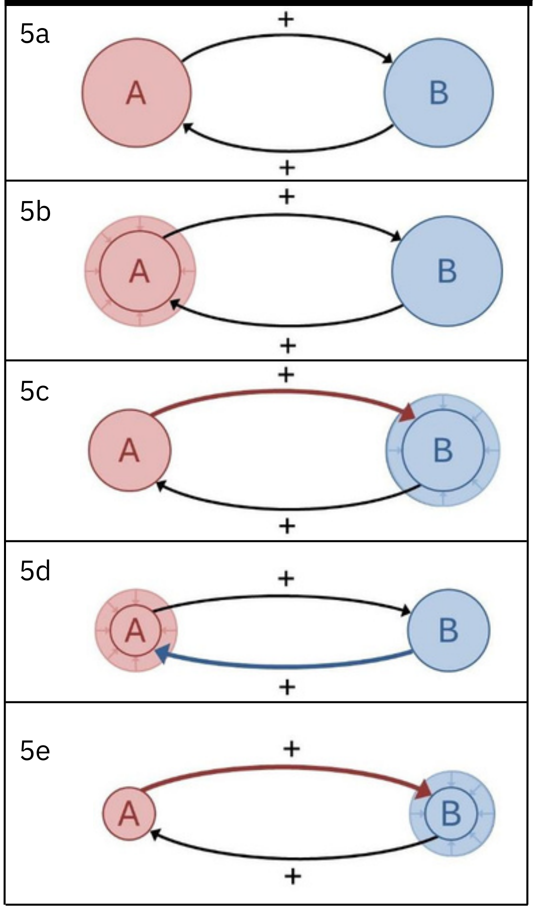

Source of common confusion 1: positive and negative causal relationships versus feedback loops

We use the terms "positive" and "negative" in reference to both causal relationships and feedback loops, so we need to remember that causal relationships and feedback loops are two different things. When they involve more than two causal factors (see the section "Multifactor feedback loops" below), positive feedback loops may be composed of both positive and negative causal relationships, as can negative feedback loops. Indeed, as shown in the example below, a positive feedback loop can involve only negative causal relationships. And negative feedback loops can be composed of nearly all positive relationships (but must have at least one negative causal relationship, for reasons explained below). So when we use the terms positive and negative, we have to be clear about whether we are referring to a single relationship or an entire feedback loop.

In the positive feedback loops presented above, both causal relationships are positive. But a positive feedback loop can also involve only negative causal relationships (Figure 6) or combinations of positive and negative causal relations (see Figure 9 below). Here, an increase in A (Figure 6b) causes B to decrease (Figure 6c). Because the causal relationship from A to B is negative, A and B change in opposite directions. The decrease in B (Figure 6c) then causes a further increase in A (Figure 6d) because the causal relationship from B to A is also negative. The increase in A (Figure 6d) then causes a further decrease in B (Figure 6e), and so on.
In this positive feedback loop, one variable rises as the other falls.

It could also go in the other direction, where an initial decrease in A causes an increase in B, which further decreases A, and so on.

Figure 6: A positive feedback loop with two negative causal relationships

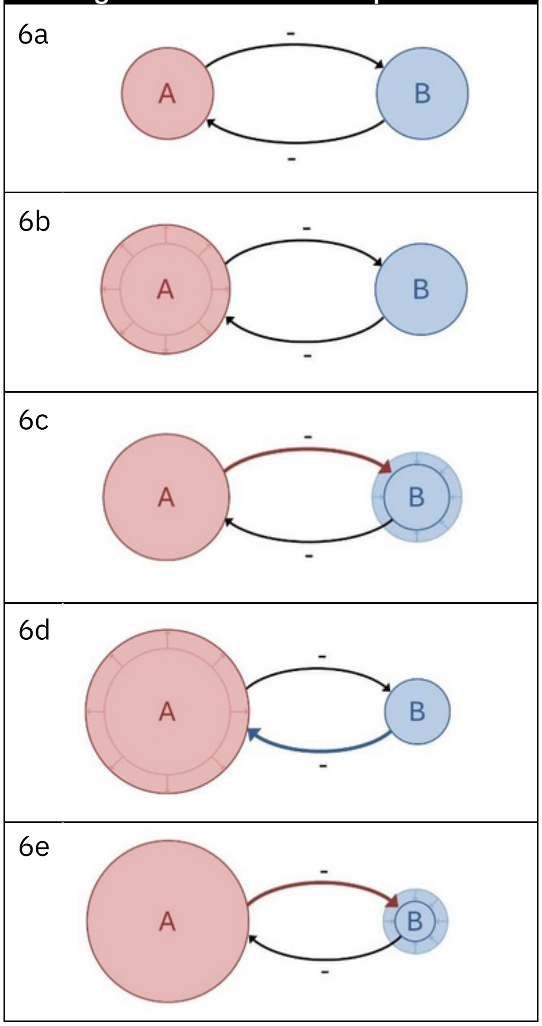

A good example of this type of positive feedback is the “Matthew effect” wherein the rich get richer by making the poor even poorer.[2] If A represents the resources held by the rich and B the resources held by the poor, then an increase in resources (Figure 6b) held by the rich allows them to pursue additional opportunities and develop new forms of exploitation, capturing (diminishing) the resources held by the poor (Figure 6c). That transfer of resources to the rich (Figure 6d) enables additional advantages with which to capture even more resources from the poor (Figure 6e). The result is a vicious cycle of growing inequality.

[2] The Matthew Effect is named for the New Testament gospel of Matthew, 25:29: “For whosoever hath, to him shall be given, and he shall have more abundance: but whosoever hath not, from him shall be taken away even that he hath.”

NEGATIVE FEEDBACK LOOPS

Figure 7: A negative feedback loop

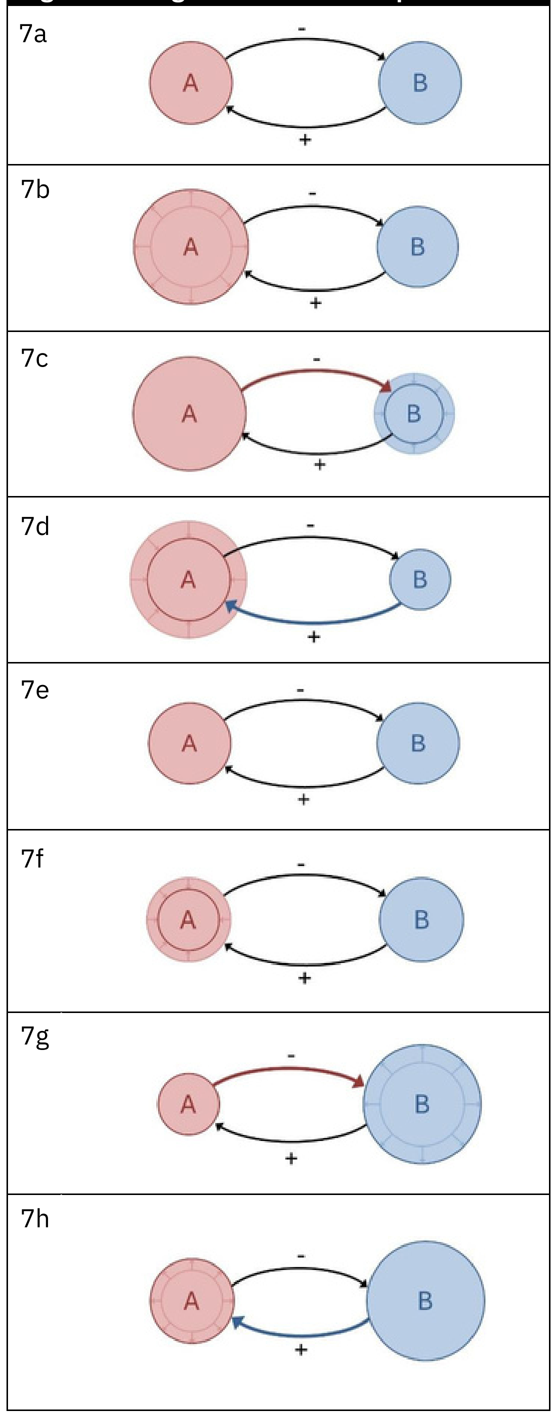

Where positive feedback loops amplify an initial change (into runaway growth, a spiral to nothing, or a paired rise and fall), **negative feedback loops** counterbalance, or cancel-out, an initial change and restore a system to an equilibrium state. They are, therefore, sometimes called "**balancing feedback loops**."

A thermostat that regulates the temperature of a room provides a good example of a negative feedback loop (Figure 7a). Let A represent the temperature of the air in the room, and B represent the temperature of air coming out of a vent (from either a furnace or air conditioner) into the room. If the room temperature rises (A in Figure 7b), the thermostat triggers the air conditioner to lower the temperature of air coming out of the vent (B in Figure 7c). This is a negative relationship from A to B; the two temperatures move in opposite directions. The drop in air temperature from the vent then causes the room temperature to decrease as well (Figure 7d). This is a positive relationship from B to A; the two temperatures move in the same direction. The room returns to its set temperature, the thermostat stops the air conditioner, and the vent temperature also returns to room temperature (Figure 7e).

If the room temperature drops (A in Figure 7f), the thermostat triggers the furnace to increase the temperature of vented air (B in figure 7g).

The extra heat from the vent then raises the temperature of the room (Figure 7h). And the thermostat turns the furnace off so that both the room air and vent air temperatures return to their initial, programmed setting (Figure 7a).

Balancing feedback loops restore a variable to a certain value after a disturbance, and thereby maintain consistent conditions.

In the example above, the thermostat cancels out a change to the set temperature then switches off the causal connections (the air conditioner or furnace) until another change in room temperature occurs. Other negative feedback loops produce continuing oscillations instead of the kind of static equilibrium we see in the thermostat example.

The fluctuations between predator and prey populations in an ecosystem (Figure 8a) provide a good illustration of an **oscillating negative feedback**. Imagine A represents the population of a prey species (such as rabbits) and B represents the population of a predator species (such as wolves). An increase in the prey population (Figure 8b) provides more food to the predator population, enabling it to breed and grow in numbers (Figure 8c). But the larger predator population over-consumes (reduces) the prey population (Figure 8d), and the decrease in food then causes part of the predator population to die off (Figure 8e). The decline of the predator population enables more prey to grow (figure 8f) and the cycle repeats.

The result is a fluctuating pattern where one variable is rising as the other is falling, and vice versa. This represents a dynamic equilibrium, because it is a persistent pattern but does not settle at one value (such as the temperature setting of the thermostat).

Figure 8: An oscillating negative feedback loop

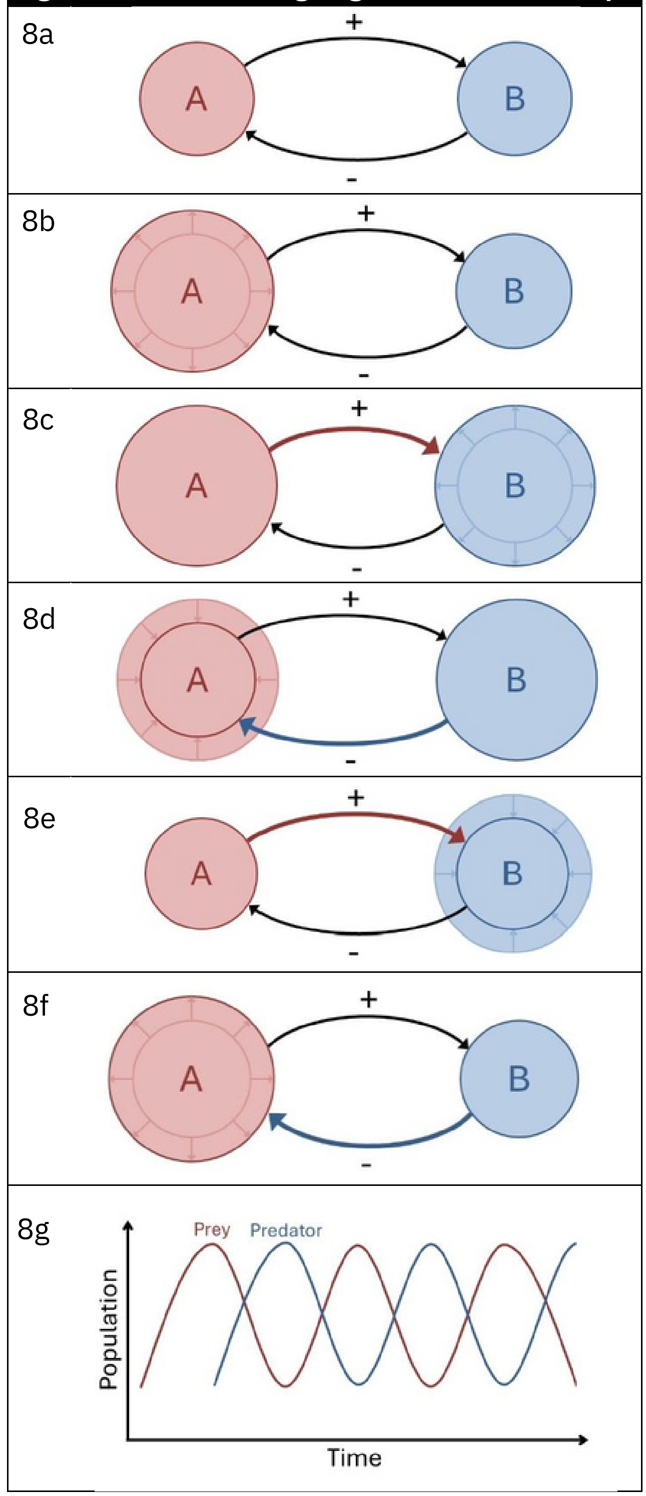

Source of common confusion 2: positive and negative as empirical vs. normative descriptors[3]

In everyday conversation, we use the terms “positive” and “negative” to mean good and bad, respectively. In regard to the relationships and feedback loops of CLDs, however, the terms “positive” and “negative” do not indicate such value judgements. The terms are used in an empirical rather than normative manner, concerning “what is” rather than “what ought to be.” The terms simply indicate whether relationships go in the same direction (positive) or opposite directions (negative), and whether a feedback loop is self-reinforcing (positive) or self-balancing (negative). Positive feedbacks can be desirable or undesirable, as can negative feedbacks; it depends on the system in question. Desirable positive feedback loops are often called “virtuous cycles” and harmful positive feedback loops are often called “vicious cycles,” but in both cases the feedback loop is positive (self-reinforcing). The table below emphasizes this distinction by providing additional examples.

Type of Feedback	Impacts of Feedback: Beneficial	Impacts of Feedback: Harmful
Positive (self-reinforcing)	**A desirable change in one thing produces a desirable change in second thing, which increases the desirable change in the first thing, which increases the desirable change in the second, and so on. We saw this scenario in the example of compound interest, where our bank account balance and the interest we earn amplify one another.**	**A harmful change in one thing produces a harmful change in second thing, which increases the harmful change in the first thing, which increases the harmful change in the second, and so on. We saw this scenario in the example of the Matthew effect in which the rich get richer by making the poor poorer.**
Negative (self-balancing)	**A shift away from a desirable state triggers counter-acting forces that restore the desirable state. We saw this scenario in the example of the thermostat that maintains comfortable living conditions by counteracting rises and falls of temperature.**	**Shifts away from an undesirable status quo trigger counter-acting forces that restore the undesirable status quo. Rebellions against government repression, for example, may temporarily increase popular rights and freedoms, but often trigger government crackdowns that reverse these gains.**

[3] The terms “positive” and “negative” can be confusing in multiple ways, so some authors use alternative terminology. The terms “positive” and “negative,” however, are so well established and commonly used that you must understand them if you are to accurately interpret and describe many causal loop diagrams.

MULTIFACTOR FEEDBACK LOOPS

The feedback loops examined so far have involved only two elements, but feedback loops often involve many causal factors in a long chain of cause and effect that cycles back to the element that initially sets change in motion. There must be such a cycle for there to be a feedback loop.

Figure 9a presents a four-factor feedback loop. To determine whether it is a positive feedback loop or a negative feedback loop, we can imagine an initial increase in causal factor A and think through the changes that result for the other three factors (Figure 9b, moving clockwise through the relationships).

Figure 9a: Feedback loop with four causal factors

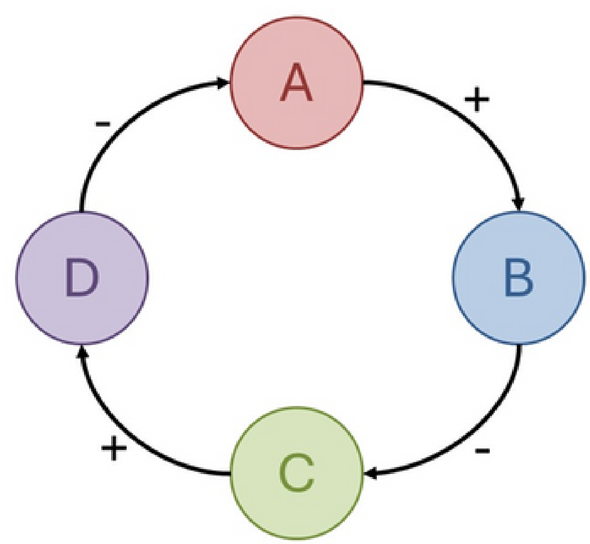

Figure 9b: Feedback loop with four causal factors and an initial increase

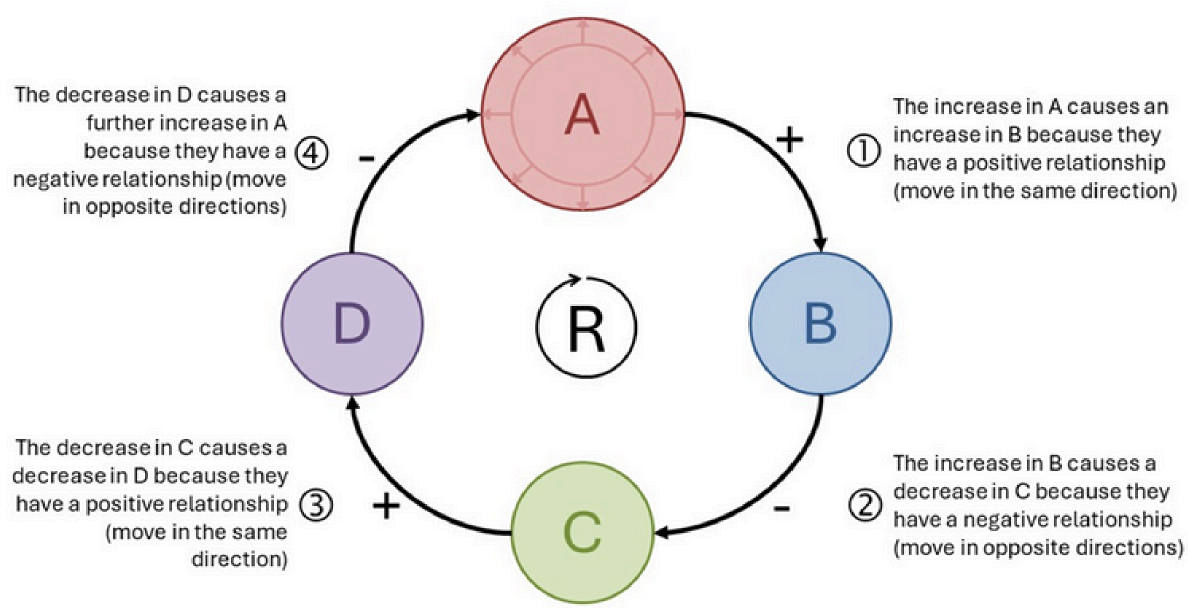

In Figure 9b, the initial increase in A sets off a series of changes that amplify the change in A. This is a positive (self-reinforcing) feedback loop. We get the same amplification of the initial change if we start by decreasing (rather than increasing) A (Figure 9c). Note that Figures 9b and 9c, unlike the ones above, do not graphically depict the increase or decrease of each element (only the initial change). These Figures instead describe such relationships textually. Following a commonly used convention, the "R" symbol in the middle labels the feedback as a reinforcing (positive) one.

Figure 9c: Feedback loop with four causal factors and an initial decrease

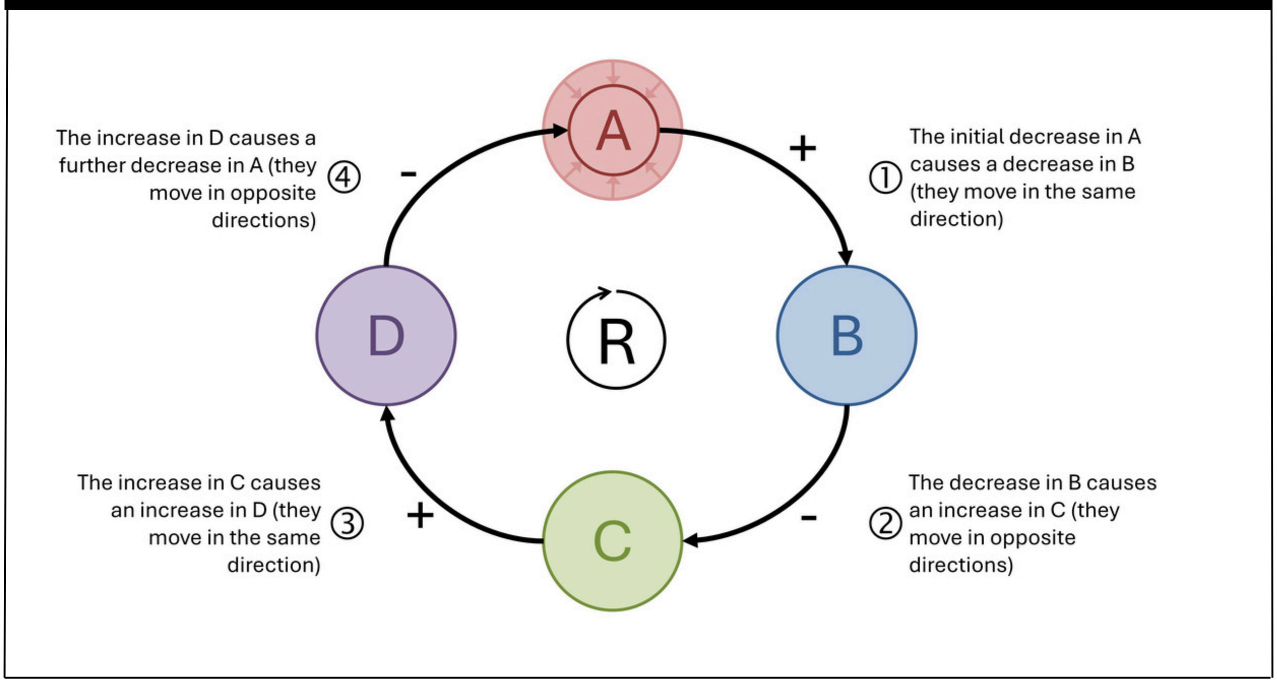

Figure 10 presents a five-factor feedback loop. Is it a positive or negative one? We can think it through step by step like we did with Figure 9, or we can take a shortcut. A quicker way to determine whether a feedback loop is positive or negative is to multiply the plus and minus signs one by one through the loop to determine whether the product—and the feedback loop—is positive or negative.

$$(-) \times (+) \times (+) \times (-) \times (-) = (-)$$

Figure 10: A feedback loop with five factors

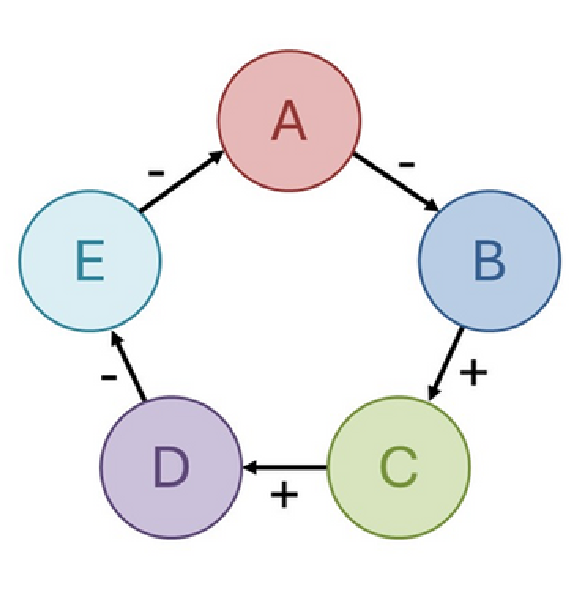

A negative times a positive is a negative, times a positive is negative, times a negative is positive, times a negative is negative. This is a negative feedback loop. Whether A increases or decrease, it sets off a chain of cause and effect that counteract the initial change.

Alternatively, we can count the number of minus signs in the feedback loop. If there is an odd number of minus signs, it is a negative feedback loop. If there is an even number of minus signs, or no minus signs, then it is a positive feedback loop.

STEP-BY-STEP INSTRUCTIONS

The process of creating causal loop diagrams involves three broad stages that often overlap in practice. In each stage, we present tips derived from our own experience that may help improve your diagrams.[4]

1. **Identify the relevant elements (causal factors) and connections (causal relationships).**

A CLD that includes every element and connection in a system will be utterly illegible and practically useless. A helpful CLD will focus instead on the causal factors and relationships that are most influential upon the system behaviour or aspect of the system you are investigating, acknowledging that the diagram is a simplified representation of a much more complex reality. We have found that CLDs with more than twelve elements tend to overwhelm viewers and lead to more confusion than understanding.

Here are five tips to help you most constructively identify and label system elements:

- Recall from the Elements section above that an element can be any causal factor that shapes system behaviour, provided that it can increase and decrease. A factor that cannot increase or decrease—for example, a particular actor—will not work in a CLD but may be reformulated in a way that does—as a particular actor's power, for example, which can rise and fall.
- Each element should be distinct (separate, non-redundant, and logically independent) from the other elements—that is, not just a slightly different version of another element. If a CLD of a geopolitical system includes the elements "economic power" and "gross domestic product," for example, we might find these causal factors to be redundant and difficult to distinguish. It would be better to collapse them into a single element.
- Elements should be labelled as *things* rather than the *absence* of things. For example, use the variable "polarization" rather than "lack of political unity," and "trust" rather than "absence of trust."
- Do not label your variables with adjectives that indicate increase or decrease. For example, use "cooperation" rather than "rising cooperation" or "declining cooperation." This and the previous convention allow the arrows to indicate the directionality of change rather than building it into the elements.
- Label elements as concisely and specifically as possible to make your diagram easily legible. Instead of "possibility a country will launch a first strike using nuclear weapons," for example, use the label "risk of nuclear warfare." You can elaborate and clarify your labels in the narrative that accompanies the diagram.

[4] This section is indebted to Barbrook-Johnson and Penn (2022, pp. 51-54), who set out seven steps for creating causal loop diagrams.

To identify the (twelve-or-so) most important causal factors and the relationships between them, you can draw upon information about the system from a variety of sources, including surveys, expert elicitation, literature scans, intuitions, participatory discussions with stakeholders, datasets, and quantitative or qualitative studies. In this first stage, your goal is to produce an inventory of relevant causal factors and causal relationships to be included in your CLD.

2) **Lay out the elements and draw the (positive and negative) causal relationships (arrows) between them.**

Where the first stage is about consolidating your knowledge of the system, the second stage is about drawing your CLD in a way that most clearly and accurately represents that knowledge. You can draw your CLD by hand (which may be easiest way to produce the first few—inevitably messy—drafts) or with a software. We produced all the diagrams in this handbook using Microsoft PowerPoint, but you may want to explore more specialized software, such as Vensim or Kumu.

To the extent that elements can be prioritized, start by laying out those at the top of your inventory, draw the connections between them, then rearrange the elements and connections to be as clearly organized as possible. Add additional elements and connections and continue to reorganize them as necessary. A viewer friendly CLD will avoid criss-crossing arrows, which will likely require you to reconfigure your layout several times, and some intersections may be unavoidable. Feedback loops should be readily visible and labelled as positive or negative (R or B, respectively). Note, however, that not all causal factors must be part of a feedback loop; some may simply affect, or be affected by, other factors that are part of a loop (see example 3 below).

With each draft of your CLD, you should go back to your information sources and solicit others' reactions to verify that your visualization captures the systemic behaviour you are investigating. Revise as needed. You may find that you have to add additional elements and connections to improve the accuracy of your CLD, or you may find that some elements and connections are not relevant to your concerns and can be removed to make the diagram more focused and concise. This second stage is an iterative process that generally involves several rounds of visual rearrangement, verification, and revision.

3) Analyze and interpret the diagram

Articulate a narrative that explains the causal relationships of your CLD in a logical, step-by-step manner. A crucial part of the analysis is to identify the feedback loops and consider whether they tend to keep a system at equilibrium or promote runaway change. Your narrative should ultimately explain how the elements, connections, and feedback loops generate the systemic behaviours you are investigating. It should consider how the system will change or resist change amidst shifting circumstances. And it should help you to develop strategies to alter or maintain system behaviours.

A helpful next step is to identify which elements are most vulnerable, and which are most influential. Vulnerable elements are those that are most highly affected by other elements—that have the most arrows pointing to them. Vulnerability may indicate that a variable is over-determined and thus especially difficult to change. Influential elements are those that most influence other elements—that have the most arrows pointing out from them. These elements may represent the most significant drivers of system behaviour and important sites for intervention.

The basic logic of CLDs suggests three broad strategies by which to intervene in the system to change its behaviour. First, you can add or strengthen an element, or remove or weaken another. Introducing the element "accountability" to a governance system by implementing the appropriate policies and institutional reforms, for example, may diminish the element "corruption." Restricting the element "backdoor lobbying" could amplify the change.

Second, you can introduce new connections between elements, or eliminate existing connections, to change the way these variables behave. Proactive vocational and retraining initiatives, for example, might build a two-way, positive connection between the elements "employment" and "green energy infrastructure" while weakening the connection between "employment" and "fossil fuel infrastructure."

And third, you can change the valence of a causal relationship from positive to negative, or negative to positive, to alter the chains of cause and effect and perhaps even flip a feedback loop from positive to negative, or negative to positive. Public opinion may rapidly flip, for example, from supporting foreign interventions to opposing them.

Once you have completed your CLD, finally, you can use it as a stepping stone to other, more formalized methods of systems analysis. By quantifying the variables and their relationships, you can translate your CLD into the mathematical models of Systems Dynamics to simulate the system's behaviours, as well as its responses to interventions.[5] And even rough qualitative and quantitative estimates of causal relationships can be entered into a Cross Impact Balance matrix to test the stability of different system scenarios.[6] But these next steps require a strong facility with causal loop diagramming, and some real-world examples might help develop this capability.

EXAMPLE 1: THE SECURITY DILEMMA

In the security dilemma, one country's decision to increase its military forces causes a second country to feel increasingly threatened. Consequently, the second country increases its military forces, and in so doing, increases the first country's perception of threat. The first country then further increases its military forces, and the cycle escalates until one event or other triggers a war that neither country actually wants.

Figure 11: The security dilemma

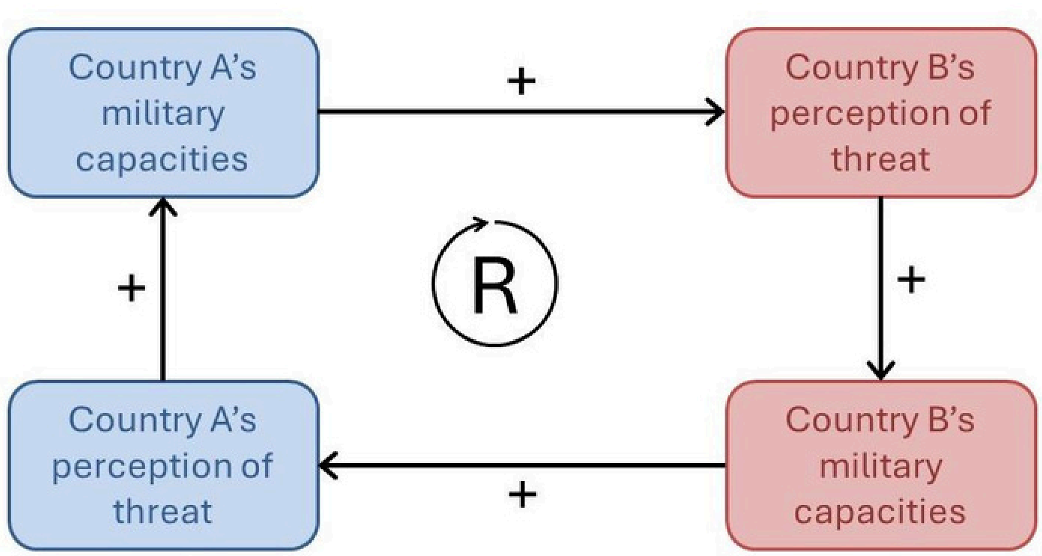

[5] For an overview of Systems Dynamics, see: Barbrook-Johnson and Penn (2022: pp. 113-138).
For an introduction to Cross Impact Balance Analysis, see: Wolfgang Weimer-Jehle (2006).

[6] "Cross-Impact Balances: A System-Theoretical Approach to Cross-Impact Analysis," *Technological Forecasting and Social Change*, vol. 73, pp. 334-361. https://doi.org/10.1016/j.techfore.2005.06.005.

This positive feedback, however, can also work in the other direction. One country may decrease its military forces to make a second country feel less threatened, so that the second country is also willing to reduce its military forces, reducing the first country's perception of threat, and so on. This positive feedback loop working in the decreasing direction captures the basic logic behind many disarmament and mutual-reassurance treaties, such as the Strategic Arms Limitation Treaties signed by the United States and the Soviet Union to reduce their nuclear arsenals (to mutually agreed levels; the positive feedback did not proceed to total disarmament).

EXAMPLE 2: THE ARCTIC ICE ALBEDO FEEDBACK

Another important real-world example of a feedback loop is the arctic ice albedo feedback. "Albedo" is a measure of how much incoming light a substance reflects, rather than absorbs. Ice, with its bright white color, has a high albedo because it reflects a high proportion of incoming sunlight back out into space. Open seawater and exposed landmasses have lesser albedos because they reflect less sunlight back out into space, and instead absorb more of that sunlight as heat.

As the Earth's temperature rises due to anthropogenic greenhouse gas emissions, arctic ice begins to melt, covering less of the Earth's surface. As the ice melts, exposed open seawater and exposed landmass both increase in area. Because water and land have lesser albedos than ice, the Earth's overall albedo – its reflectivity – decreases. As the Earth absorbs more sunlight as heat, the temperature of the climate increases, escalating the cycle.

Figure 12a: The arctic ice albedo feedback

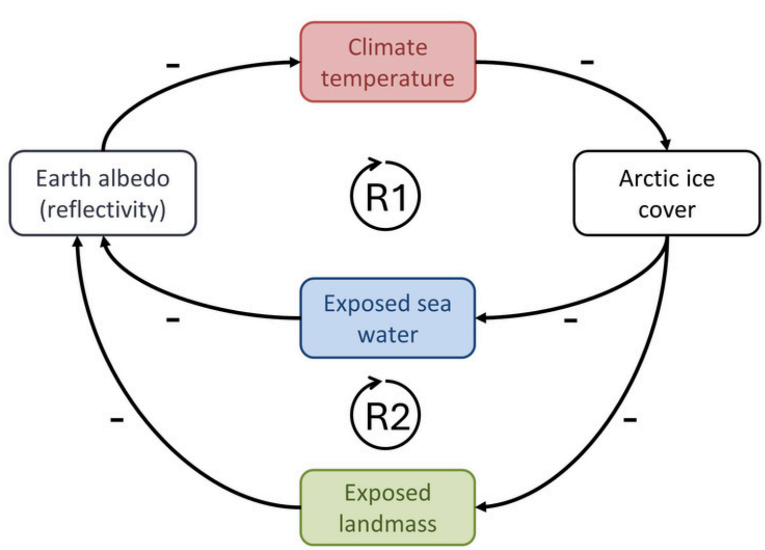

Notice that there are actually two positive feedback loops depicted in this diagram: R1, which involves increasingly exposed sea water, and R2, which involves increasingly exposed landmass. And even though all the causal relationships (signs) in this diagram are negative, the feedback loops are positive. Through these loops the earth's temperature drives its own increase, once set off by human emissions. It is a dangerous, self-reinforcing chain of events.

Some scientists hypothesize that millions of years ago, this same positive feedback worked in the other direction: declining temperatures expanded ice cover and thus reduced the area of open water and landmasses. The Earth's albedo increased, and thus caused a further decrease in planetary temperature, creating a "snowball Earth" state.

Figure 12b: The arctic ice albedo feedback with a balancing feedback loop

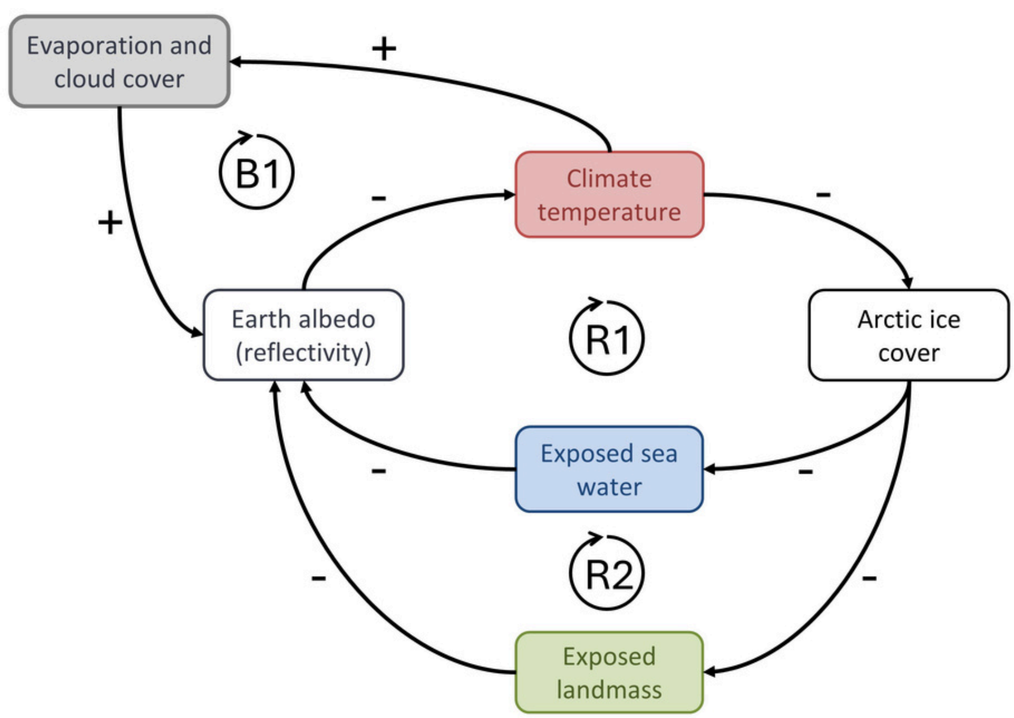

The arctic ice albedo effect also involves a negative feedback loop, labelled B1 above. As the earth's temperature increases so too does evaporation and cloud cover. Some clouds reflect more sunlight back into space (than exposed sea water and landmass), thus raising the Earth's overall albedo, and decreasing the planetary temperature. Unfortunately, this negative feedback is too small in magnitude to counteract the overall warming driven by the positive feedback loops.

EXAMPLE 3: SOLUTIONS THAT BACKFIRE

In his book *Systems Thinking for Systems Change*, systems thinker David Peter Stroh identifies a number of systems "pathologies" by which a system's structure resists efforts to change system behaviours and thereby prevents individuals and organizations from attaining their goals. One system pathology that often afflicts policymaking is what Stroh calls "solutions that backfire".

Figure 13a: Solutions that backfire

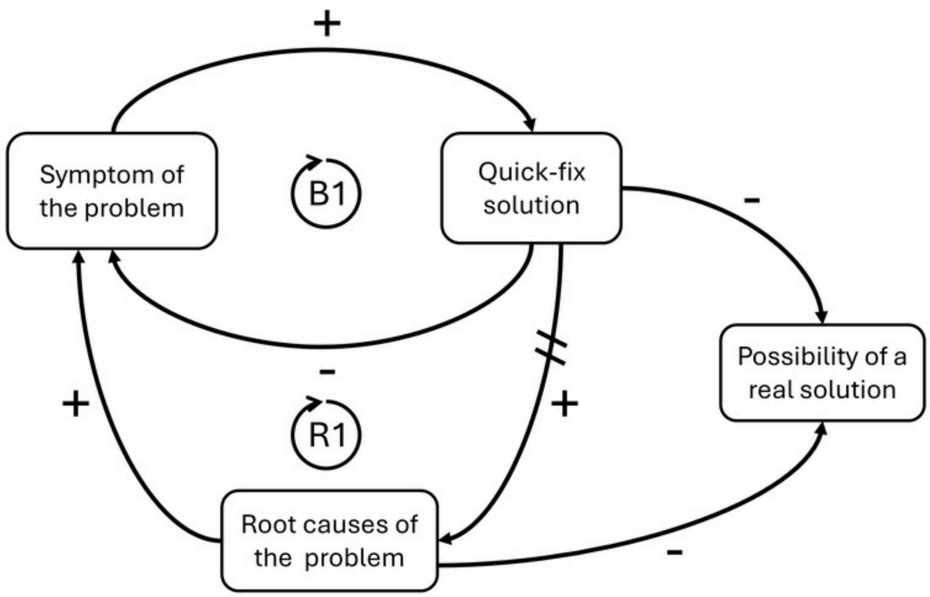

When the symptoms of a social problem grow (top-left), policymakers often pursue a quick-fix solution that reduces the symptoms. This creates a negative, self-cancelling feedback loop (B1). The symptoms may recur, but additional applications of the quick-fix solution diminish them.

Over time, however, the quick fix solution actually worsens the underlying causes of the problem (the dashes on the arrow between these two elements indicates a delay). As those root causes increase, so too do the symptoms of the problem. Policymakers continue to implement the quick-fix solution to temporarily reduce symptoms of the problem, but the problem keeps coming back even worse. This creates a positive feedback loop (R1) between the symptoms of the problem, the quick-fix solution, and the root causes of the problem.

As the root causes of the problem grow, and as policymakers become more deeply invested in the quick-fix solution, the possibilities of implementing an actual solution that addresses the root causes diminish (far right). Note, however, that the declining possibility of a real solution is not a feedback loop because it is not part of a complete causal circuit.

Tough-on-crime policies provide an unfortunately apt example of a solution that backfires. These policies respond to crime by increasing prison sentences (alongside other punitive responses). Crime may decline for a while, but tough-on-crime policies do not remedy the socio-economic factors that drive criminal behaviour. In fact, imprisonment often worsens such factors by producing even more hardened criminals who have less opportunity to reintegrate into law-abiding society. Once they are freed, many re-offend. As crime recurs, people remain ideologically committed to even harsher punishments, while the socio-economic drivers of criminality continue to worsen, so that a real solution becomes less and less likely.

A real, lasting solution to crime would reduce the root causes of criminality, as depicted in Figure 13b below. Effective policies might include social welfare provisions, community-based policing, rehabilitation programs in and out of prisons, and restorative justice measures. These policies add a new element (a different approach) to the system and may weaken the connection between “symptoms” and “quick-fix solution.” As indicated by the dashes on the arrow between real solutions and root causes, however, these policies often take time to show results. They require long-term commitments and large amounts of resources. For these reasons, real solutions may be politically unpopular and prematurely abandoned so that criminality persists, and the favored quick-fix solution continues to backfire.

Figure 13b: A real solution vs. a solution that backfires

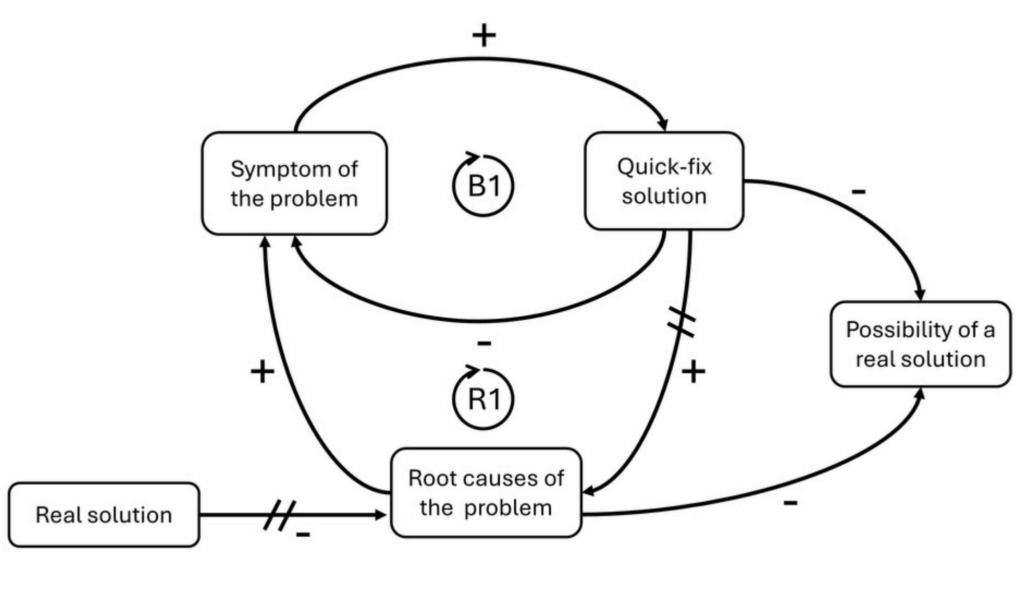

RECOMMENDED RESOURCES

- Pete Barbrook-Johnson and Alexandra S. Penn (2022). *Systems Mapping: How to Build and Use Causal Models of Systems*. Cham, Switzerland: Palgrave-Macmillan. https://doi.org/10.1007/978-3-031-01919-7.

- Donella Meadows (2008). *Thinking in Systems: A Primer*. Edited by Diana Wright. White River Junction, VT: Chelsea Green Publishing.

- David Peter Stroh (2015). *Systems Thinking for Social Change: A Practical Guide to Solving Complex Problems, Avoiding Unintended Consequences, and Achieving Lasting Results*. White River Junction, VT: Chelsea Green Publishing.

The resources below present particularly powerful and impressive uses of causal loop diagrams:

- Fanny Groundstroem and Sirkku Juhola (2021). “Using Systems Thinking and Causal Loop Diagrams to Identify Cascading Climate Change Impacts on Bioenergy Supply Systems,” *Mitigation and Adaptation Strategies for Global Change*. https://doi.org/10.1007/s11027-021-09967-0.

- Catherine E. Richards, Richard C. Lupton, and Julian M. Allwood (2021). “Re-Framing the Threat of Global Warming: An Empirical Causal Loop Diagram of Climate Change, Food Insecurity, and Societal Collapse,” *Climatic Change*, vol. 164, no. 49. https://doi.org/10.1007/s10584-021-02957-w.