Systems Thinking & Systems Modelling

A Course for Understanding Systems and Creating Systems Models

The Sustainability Laboratory

Systems Thinking and Systems Modelling Course

These materials were created as part of The Sustainability Laboratory's educational curriculum. It was developed with financial support from The Foundation for Advancement of Systems-Orientated Management Education at the University of St. Gallen, Switzerland.

The introductory module, Understanding Systems, was written by Dr. Michael Ben-Eli, the systems modelling modules were developed in collaboration with Loops Consulting and this e-book version was created by Systems Innovation.

An online version of the course is also available on the Kumu platform:
Click here to find the full online course

Introducion

By Michael Ben-Eli, founder of The Sustainability Laboratory

The sustainability challenges we face on the planet are multifaceted, complex, and interconnected. To understand them and address them effectively requires a holistic "systems" view that is still largely absent from most mainstream approaches in science, government, business, and education. It could be argued, in fact, that the many sustainability-related issues—climate change, desertification, potable water shortage, biodiversity loss, economic and social instability, and more—are the result of a non-systemic, fragmented, simplistic, and short-sighted world view that dominates our industrial civilization. Recent decades, however, have seen a significant surge of interest in holistic ways of looking at reality with the associated development of multiple frameworks and tools which, all together, have been hailed as the emergence of a new paradigm. The combined result has been referred to in the related, burgeoning literature as "systems thinking," "the system approach," or "the system view of the world."

Purpose of This Course

At The Sustainability Laboratory, we subscribe to the view that systems thinking and a deeply ingrained ability to "think systems" is an indispensable tool for future sustainability leaders. We have therefore made systems thinking and systems modelling a central part of The Lab's educational curriculum. The materials that follow constitute a key element in our Global Sustainability Fellows (GSF) program. It combines a number of teaching modules that divide logically in three essential parts: an introductory overview to systems thinking; the fundamentals of systems mapping and systems modelling; and a selection of illustrative case studies.

This short but comprehensive "course" has been designed to blend theoretical, conceptual exposure with hands-on engagement and collective building of systems models in groups. The course has been developed in collaboration with members of Loops Consulting, led by three of its leaders who, themselves, are alumni of the GSF program. The work was made possible by a grant from the Swiss, University-of-St. Gallen-based, Foundation for the Advancement of Systems-Oriented Management Education. Board members of the Foundation are, in their own right, leaders in the systems thinking movement.

Overview of the book

This course is divided into five modules.

Module One gives an introduction to the field of systems thinking, discussing its emergence as a new scientific paradigm and elaborating on some of the characteristics of a "systems" view of the world.

Module Two introduces the field of system dynamics, which is itself a methodology to help us think about issues from a systemic perspective. Here we discuss why we build models, and how they can help us to think and communicate more deeply regarding certain issues. Note that in this course we will deal only with qualitative system dynamics models, known as causal loop diagrams (CLDs), and will not be dealing with quantitative simulation models.

Module Three introduces the reader to the basic building blocks of CLDs, explaining how they use named variables and arrows to display the different cause and effect relationships that are assumed to give rise to a certain patterns of behaviour.

In Module Four, you will get the chance to apply the set of skills you acquired from the earlier modules, as we lead you through our own thought process for creating CLDs about selected issues. Here we will also introduce the concepts of stocks and flows, which are important in the field of system dynamics.

In the fifth and final module, we will discuss five case studies that demonstrate how CLDs have been used in particular situations, and the value that they provided.

The ultimate goal is that by end of this course, you should be able to understand and even build your own causal loop diagrams, as well as have a better grasp of what it means to think in systems.

Module One
Understanding Systems

YELLOW-RED-BLUE BY WASSILY KANDINSKY
IMAGE COURTESY OF GUILLAUME PIOLLE VIA WIKIMEDIA COMMONS

Module One Overview

This first module gives an introduction to the field of systems thinking, discussing its emergence as a new scientific paradigm and elaborating on some of the characteristics of a "systems" view of the world. The modules on system modelling (Modules 2, 3, 4, and 5) follow.

To access the text for Module One, please click here. When you are finished with that text, please return to this page and proceed to Module Two.

Module Two
Introduction to Systems Modelling

WASSILY KANDINSKI [PUBLIC DOMAIN], VIA WIKIMEDIA COMMONS

Module Two Overview

The purpose of this particular module is to give you a brief introduction to system dynamics, where it came from, why it is used, and what are its limitations. We will then move on to explore the actual building of system dynamics models in the third module.

By the end of this module, you should understand:

What mental models are
Why we use system dynamics
The context in which system dynamics arose
Some of the typical ways in which system dynamics is used
Some of the benefits of using system dynamics
Some of the central assumptions of system dynamics
Some of the limitations of systems dynamics

This module is split up into five sections, described below:

1. Mental Models

2. An Overview of System Dynamics

3. A Brief History of System Dynamics

4. The Central Assumptions of System Dynamics

5. The Limitations of System Dynamics

Section 2.1: Mental Models

Let's begin this section with a few questions to consider:

Will the increased construction of homeless shelters solve the problem of homelessness in Dublin, Ireland?

How might increasing access to electricity in Tanzania affect educational outcomes?

How might increasing rooftop solar installations in the US affect the price of electricity from the grid?

How might lowering working hours affect the material wealth of a country?

Would universal basic income bring about more or less wellbeing in a nation?

Are global warming and the war in Syria in any way related?

How might the increasing use of biofuels affect global food prices?

Can the cutting of trees in a forest reduce that forest's vulnerability to wildfires?

How does your perception of other people affect your behaviour towards them? And how does your behaviour towards them affect their perception of you?

Exercise

Before we go further, we ask you to spend some time thinking about your own responses to these questions (the list of questions can be found again below). We recognise that these questions are quite complex and that you are unlikely to have an answer ready at hand. Nonetheless, we ask that you make some attempt at answering them, in whatever way you see fit.

We strongly encourage you to write each answer down on a separate piece of paper, leaving one side of the page blank. We will ask you to come back to these answers towards the end of the course to see if your thinking or the way you approach your thinking has changed in any way based on what you have learned.

You don't need to answer every question, just the ones that are most interesting to you. We'd ask that you spend no more than 10 minutes in total answering at least two of these questions.

You are likely to have to make many assumptions in your answers, some of which could be supported by research (if you feel inclined!) and some of which can be based only on your own intuition/common sense. We encourage you to explain your answers in as much detail as you can. However, keep in mind that we are not asking you to give the “right” answers to these questions (perhaps there is no such thing), but just to use this exercise as a way of encountering the challenges of answering complex questions and explaining the logic behind your answer.

Feel free to make your answers specific to your own context. For example, instead of thinking of Dublin, Ireland in Question 1, you can think about your own city: i.e., Will homeless shelters solve the problem of homelessness in my city or hometown?

Mental Models

In writing your answers to these questions, you are likely to be using what systems thinkers would call your "mental model" about the problem or "system" at hand. In this course we will use the term mental model as follows:

Mental models are deeply ingrained assumptions, generalizations, or even pictures of how things work that influence how we understand the world and how we take action. (Taken from Peter Senge's 1990 book, The Fifth Discipline)

For example, in response to the previous question about how lowering working hours might affect the material wealth of a country, you might have said that lowering hours would reduce a country's material wealth, and that opinion might be based on a certain mental model, a certain set of assumptions about how the world works. For example, you might have argued as follows:

“I think that lowering working hours in a country will result in less material wealth for that country's population. This is because lowering working hours will mean less time is spent producing goods and providing services, which results in less income generated per person."

Or you might hold a different mental model about the situation, which would lead to a different answer. For example, you might have argued as follows:

“I think that lowering working hours will result in greater wealth for the population because the working population will have more time to rest, which makes them more productive in their working hours. This increased productivity will more than make up for the reduced hours such that there will be more goods and services produced and therefore, more wealth to go around.”

Either of these answers and corresponding mental models might be right (or you might consider both to be wrong) and either might lead to different actions. For example, a person with the first mental model might vote against legislation for lowering working hours, whilst a person with the second mental model might vote for it. In general, our mental models about how the world works tend to inform not only the opinions we hold, but also the policies we vote for, the decisions we make, and even the way we behave.

What the previous page intends to show is that while most of us do not always build formal models of how we think, we are still relying on models all the time. The only difference is that these models are mental models, something that we keep 'in our head,' so to speak. They are rarely something that we express to others or to ourselves in significant detail.

For example, when we answer a question such as those presented at the start of this module, we rarely make all/most of the assumptions behind our arguments crystal clear to others, and even to ourselves. This is because we often hold in our minds such a complex web of assumptions that it is hard to explain them all by words alone. In many cases we may not even be aware of this web of assumptions, because we don't take the time to reflect on them.

Additionally, for many topics we may not have formed any mental models at all. Some questions/topics may seem too complex for us to deal with, as there are too many things to consider all at once. Alternatively, we might naively presume that we know the answer without thinking about it in the depth required.

Summary

That bring us to the end of section 2.1. We hope that by now you understand what we mean by the term mental model, and can reflect on many of your own mental models about the world. In the next section we will introduce the methodology of system dynamics.

If you need to return to the table of contents, please click here. We will include these links at the end of every section to make course navigation a bit easier.

Section 2.2: System Dynamics

System dynamics is a methodology for building formal models of systems, and there are many different types of systems that system dynamics deals with: anything from an economy to an ecosystem to a certain part of a healthcare system. Rather than keeping our mental models about such systems "in our head", system dynamics offers a way to formally map out those mental models. In doing so, it also offers an opportunity for us to reflect on our mental models about certain systems, and to perhaps alter them when we see fit, i.e. when we see a discrepancy between our previous understanding of how a system works, and how it appears to work based on our new analysis. This allows us to make better informed decisions that are (hopefully) based on more accurate models about the world.

In this course, we will be focussing on a specific segment of system dynamics, known as causal loop diagramming. Causal loop diagramming is essentially the qualitative version of system dynamics modelling, as opposed to the quantitative aspect that deals with simulation models, i.e. quantified models that use equations to simulate how a system might behave over time under a given set of assumptions and conditions.

Causal loop diagrams (CLDs) are not quantified. Instead, they are qualitative diagrams that map out the cause and effect relationships that we assume give rise to a certain problem/pattern of behaviour. They help us to analyse why certain problems may be occurring, and why certain solutions may or may not work.

We will begin demonstrating what CLDs look like in the third module!

Complex Systems

System dynamics is typically only used when one is dealing with a complex system. A system is a very broadly applicable term that can be defined as a number of parts that are interacting, often with the aim of achieving a certain goal. For example, the human body is a system made up of your heart, your lungs, your veins, your brain, and much more. All of these parts are constantly interacting, and together they provide you with the incredible gift of life. On a more abstract level, a school is also a system, made up of teachers, buildings, students, books, etc., all working together towards the purpose of providing education.

A system can be complex when there is a a lot of interaction between its parts. This makes it hard to predict how affecting one part of the system will impact the rest. Note that systems can be complex even if there are only a few variables: it is the level of interaction between variables, rather than the number of variables, which makes a system complex. Nonetheless, more variables in a system can often lead to more interaction, and thus more complexity.

Examples of complex systems include ecosystems, economies, planet Earth, social networks, and much more.

Unintended Consequences

The high degree of interdependency in complex system can make it extremely hard to predict how affecting one part of the system will affect the rest. Certain actions can ripple along unacknowledged chains of cause and effect, resulting in consequences that are neither intended nor desired. This explains why many "solutions" to problems in complex systems may end up being ineffective or, worse still, damaging to the very problem they are trying to solve.

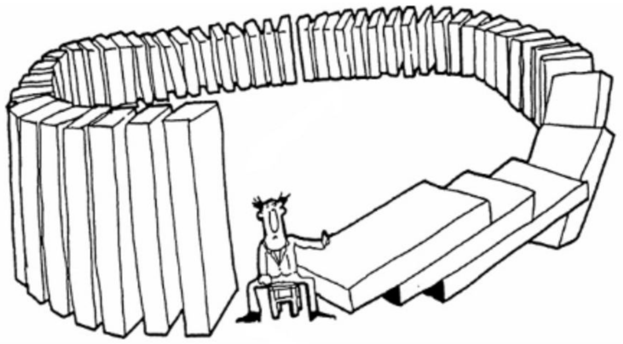

For example, in the extremely complex system of the global economy, it is often very difficult to predict how a certain policy will play out. Subsidizing biofuels may go some way to alleviating the problem of greenhouse gas emissions, but will the increased demand for crops also put pressure on arable land and increase the price of food, with potentially disastrous effects?

Another example is ecosystem change, where it is very hard to tell how affecting one species in the system will affect the rest of the system, given the high level of interdependency between species in the same ecosystem. Consider how bees play a crucial role in pollination, and how pollination plays a crucial role in providing future food for bees and many other species in the same ecosystem.

As such, we see that when a problem arises in a complex system, and some kind of intervention seems required, it is extremely important that we try to develop an understanding of the system before intervening in it. In the same way that an investor will perform “due diligence” on a market before investing in it, we must perform our own due diligence on a system before intervening in it. System dynamics helps us to perform this due diligence, so that we find better solutions to problems in complex systems, and avoid - or at least become aware of - any unintended consequences of our strategies.

Delays, Feedbacks, and Nonlinearities

Apart from a high level of interdependence, systems can also be complex due to the existence of delays, feedback loops, and nonlinearities. Here we will discuss each of these system characteristics, and why they make a system complex.

Delays: When there is a significant delay between a cause and an effect, it can be hard to see the relationship between the two, and to predict how affecting one will eventually affect the other. Consider how a firm might make orders for new inventory based on existing stock, but that those orders won't arrive for some time (there is a delay between making the order and the order arriving), by which time the inventory may be smaller or larger than expected. This can make the managing of stock a complex affair.

Feedback Loops: feedback loops refer to circular chains of cause and effect. For example, when an increase in population leads to an increase in births, which leads to a further increase in population (this is known as a reinforcing feedback loop, a topic discussed in more detail in the next module).

Nonlinearities: these are relationships whereby an increase in the cause variable does not produce a consistent change in the effect variable. As an example, consider how an increase in the price of a good from $1 to $2 might have a much bigger effect on demand than an increase from $3 to $4.

System dynamics has ways for explicitly highlighting the existence of delays, feedback loops, and nonlinearities in a system, and this helps one to recognise the complexity of that system, and find solutions that account for this complexity. Causal loop diagramming is particularly useful for highlighting feedback loops in a system, and this is something we will be dealing with extensively in this course.

Group Model Building

Causal loop diagramming is used not only to model your own thoughts, but also the thoughts of others and of groups that we are a part of. In what is known as Group Model Building, different stakeholders to a problem gather together and, in a facilitated process, create their own shared system dynamics model (which can be qualitative or quantitative, or both). By doing so, a shared understanding of a problem can be developed, power can be distributed between different stakeholders, and ownership of the solutions identified can be created. This method serves as a great opportunity for increasing stakeholder participation in decision making, something which is becoming increasingly recognised as important, particularly within the sustainability realm.

We will see examples of this method being used later in the course!

The Benefits of System Dynamics

To summarise, below you can see a list of some of the benefits that can arise from developing a system dynamics model. The model and the process of building it can help us to:

Think deeply and logically about complex questions/problems, such as those posed at the beginning of this module
Explain our reasoning/our mental models to others

Make the assumptions in our thinking more explicit to others and even to ourselves

Develop more accurate mental models, i.e. ones that better reflect reality

Forecast different scenarios and identify leverage points that can lead to effective and long-term solutions

Develop surprising and often counter-intuitive insights about problems, which leads to innovative solutions

Think about potential unintended consequences of proposed solutions

Sections Summary

That brings us to the end of section 2. We hope that by now you are aware of what system dynamics is, as well as why and in what context it is used.

In the next section we will take a few moments to understand the context in which system dynamics arose. This will also give you a feel for the various ways in which it is applied.

Section 2.3: A Brief History of System Dynamics

System dynamics was developed in the 1950s by a man called Jay Forrester. Forrester originally had a background in engineering and was one of the early innovators in computer science. Despite his success in these fields, he decided to switch his focus to management as, in true systems fashion, he felt that what he'd learned in engineering systems could be also useful in human and social systems. He also felt that management was the realm where the most pressing problems of the time lay.

Forrester had the opportunity to explore his potential contributions to the field of management when he was offered a professorship in the newly-formed MIT Sloan School of Management. His initial goal was to determine how his background in science and engineering could be brought to bear, in some useful way, on the core issues that determine the success or failure of corporations.

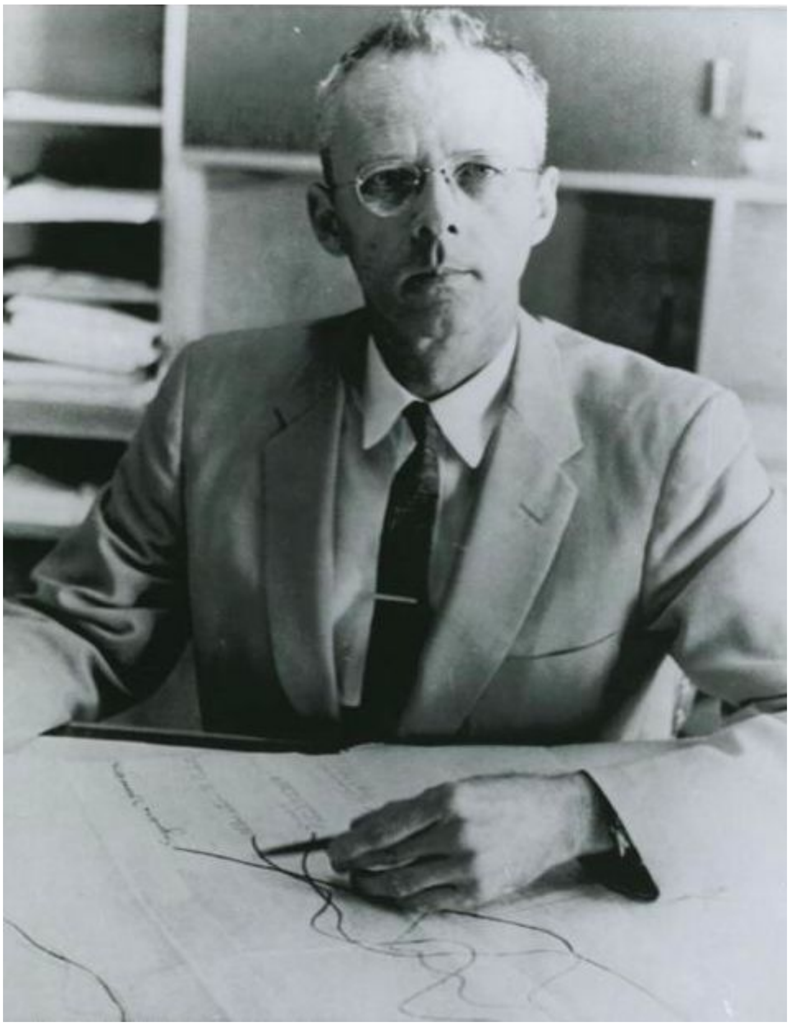

Jay W. Forrester, August 22 1964; image courtesy of MIT Museum

Industrial Dynamics

Forrester's first experience of the value of systems modelling came from conversations he had with managers at General Electric. These managers were confused about why the employment levels in their household appliance plants exhibited three year cycles, whereby they would have to lay off a significant number of staff every three years. It was easy to say that business cycles caused fluctuating demand, but that did not seem like a sufficient explanation to Forrester.

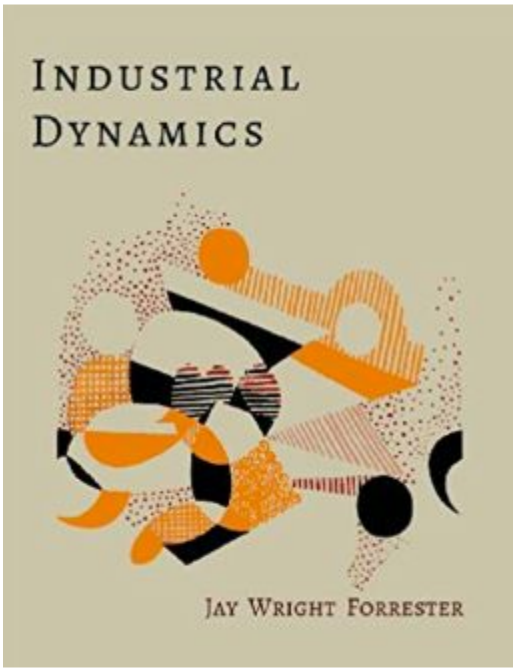

He began to ask about how the corporation made hiring and inventory decisions, and using a pen and paper he started to draw out the cause and effect relationships that represented the organisation's hiring policies, and how they related to inventories and orders. From this model, he could 'simulate' (by doing his own equations on pen and paper) how many people would be hired on a week by week basis. This first pencil and paper inventory control system was the beginning of system dynamics. It was described in Forrester's first publication to use the systems dynamics methodology, which was titled Industrial Dynamics.

Forrester's analysis revealed that even if the firms incoming orders remained constant (i.e. even if there was no business cycle), employment instability could still arise as a consequence of common decision-making policies. This highlighted system dynamics' ability to reveal unintended consequences of policies in complex systems, a theme we discussed a few pages ago.

Urban Dynamics

In the years that followed, Forrester and a team of graduate students moved the emerging field of system dynamics forward at a rapid pace, and specific software was developed for the method, allowing models to be simulated more easily than on pen and paper as Forrester had originally done.

From the late 1950s to the late 1960s, system dynamics was applied almost exclusively to corporate/managerial problems. In 1968, however, Forrester began conversing with John Collins, the former mayor of Boston, who happened to be in the office next to him. From these conversations, Forrester developed the work known as Urban Dynamics, which was the first non-corporate application of the system dynamics methodology.

The work was, and is, very controversial, because it shows why many well-known urban policies are either ineffective or make urban problems worse. Further, the model shows that counter-intuitive policies - i.e., policies that appear at first glance to be incorrect, often yield surprisingly effective results. As an example, in the Urban Dynamics model, a policy of building low income housing creates a poverty trap because such housing used up space where jobs could have been created, while drawing in people who need jobs. Forrester's analysis, aided by his model, said that building low-cost housing was a powerful process for producing poverty, not alleviating it.

This case again demonstrated how system dynamics could be used to discover the unintended consequences of certain actions or policies, and it also showed how these models can help generate counterintuitive insights that lead to the discovery of new and unexpected solutions to big problems.

World Dynamics and the Limits to Growth

Shortly after Urban Dynamics, Forrester was asked by the Club of Rome to use his new methodology to aid a study on what the Club called the "predicament of mankind": the global crisis that may appear some time in the future, due to the demands being placed on Earth's carrying capacity (i.e. its sources of renewable and nonrenewable resources and its sinks for the disposal of pollutants) by the world's exponentially growing population.

In response to this request, Forrester created a system dynamics model of the world's socioeconomic system, naming it WORLD1, which was later refined to WORLD2. The model mapped important interrelationships between world population, industrial production, pollution, resources, and food. It showed a collapse of the world socioeconomic system sometime during the twenty-first century, if steps were not taken to lessen the demands on

Earth's carrying capacity. The model was also used to identify policy changes capable of moving the global system to a fairly high-quality state that is sustainable far into the future.

A further iteration of the model, known as WORLD3, was the basis of the famous book, The Limits to Growth, in which some of Forrester's students delivered his insights in a more reader friendly way. Both World Dynamics and Limits to Growth received strong public reactions, and the discussions they sparked were an important step in increasing awareness about the necessity of addressing the sustainability challenges facing the planet. The link between system dynamics and sustainability has been strong ever since.

System Dynamics Today

Before he passed away in 2016, at the age of 98, Forrester spent his remaining years working on two main things: building a system dynamics model of the US economy, and spreading the skill of systems thinking and systems modelling in education, from kindergarten right through to university. Alongside him, countless practitioners have taken the field further and applied system dynamics in many different industries, including business strategy, public policy, finance, healthcare reform, environmental management, and most recently it has even been applied in the modelling of diseases within the human body. Many of the results have been truly fruitful,

Section Summary

Having given some context to the development of system dynamics, it will be worthwhile to spend some time thinking about some of its foundational assumptions, many of which are common to the field of systems thinking in general. This will be the focus of the next section.

Section 2.4: Central Assumptions of System Dynamics

In the pages that follow, we will discuss two of the core tenets upon which system dynamics is based. These are that structure drives behaviour and that it is better to adapt the endogenous perspective when solving problems.

Structure Drives Behaviour

We will first discuss the central belief that a system's structure is what drives its behaviour. To better understand what this means, we can use the iceberg analogy. Systems thinkers use the iceberg analogy to show that we can look at issues at four different levels: the event level, the level of patterns, the level of system structures, and at the level of mental models.

The Level of Events

Seeing the world as a series of events is like seeing the tip of the iceberg: you are only seeing the surface.

For example, imagine that you get a cold. You might react to this event by taking some medicine to treat the cold. Reacting to events is often necessary, but in the long run it may be necessary to look deeper to understand why these events are happening in the first place.

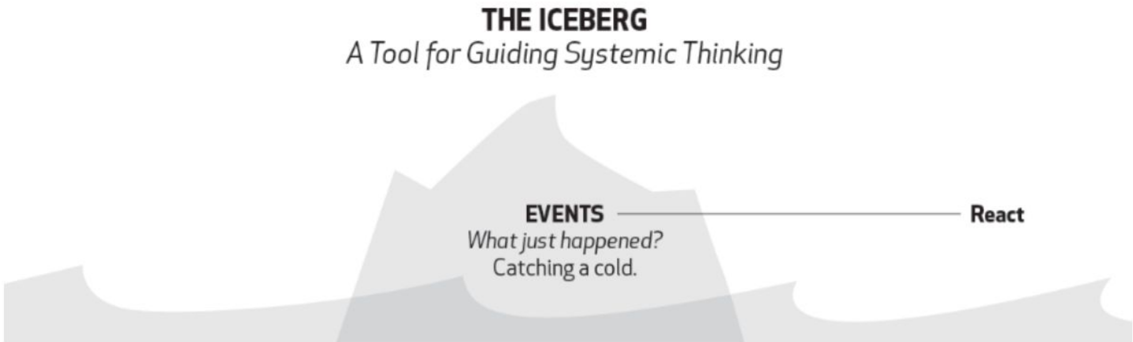

Image courtesy of nwei.org/iceberg/

The Level of Patterns

That brings us to the level of patterns, which requires us to look a little below the surface of everyday events. For example, we might recognise that we are getting a cold every few months, usually during times when we aren't sleeping as much.

Seeing patterns can help us to anticipate events, which is slightly more effective than just reacting to them.

The Level of System Structure

If we look deeper still, we will find what systems thinkers call the system structure. This is the set of cause and effect relationships that produce a certain pattern of behaviour. This does not only relate to physical structures (e.g. the human body) but also to structures in organisations (e.g. hierarchies), policies (e.g. a country's constitution), and even cultures (e.g. the relative importance that a society places on work over leisure).

In the example of getting a cold every few months, you might say that the system that is responsible for this is your lifestyle. Your lifestyle is essentially determined by the relationships between things like your work, your leisure, your income, your diet, and your health. Different relationships between these entities will produce different patterns: a lifestyle in which you spend more time at leisure and more income on healthy food will likely result in less colds than a lifestyle that spends more time at work and less income on improving/maintaining your health.

In system dynamics, the goal of modelling is to reveal the system structure as it is, so that we can understand why certain patterns and events are occurring.

The Level of Mental Models

If we go deeper still we will arrive at the ultimate cause of the system structure and the patterns and events we see: this is the level of mental models. As stated previously, mental models are deeply ingrained assumptions, generalizations, or even pictures of how things work that influence how we understand the world and how we take action.

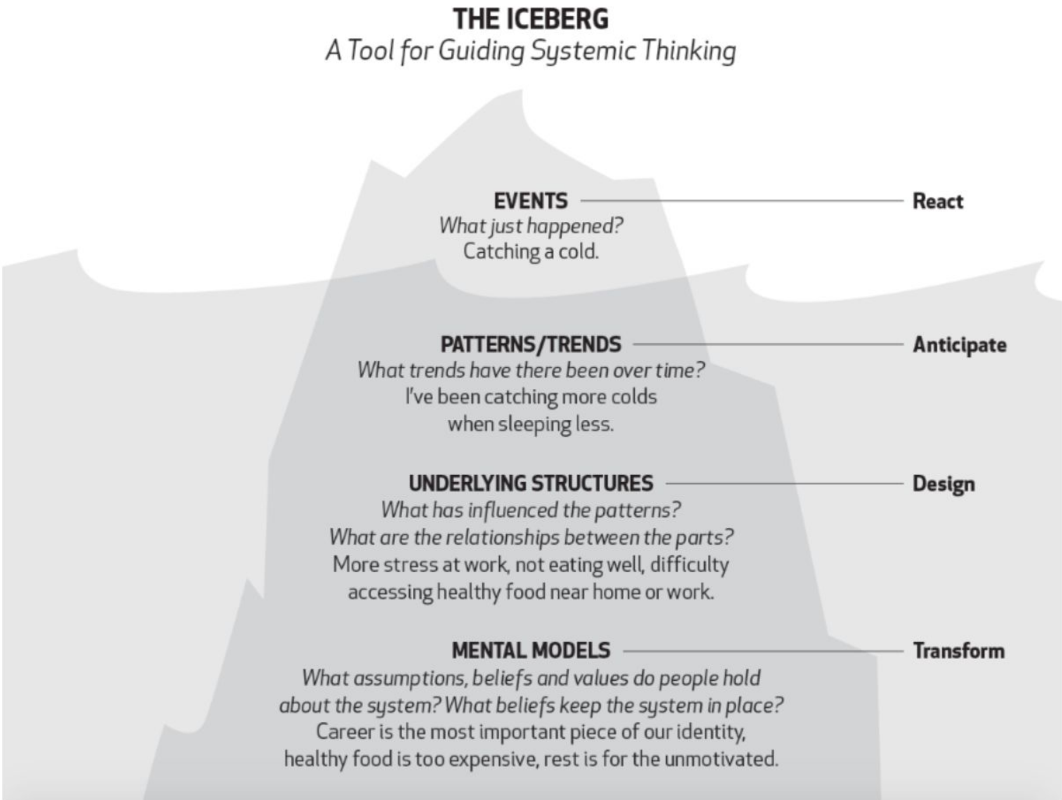

In the lifestyle example, we can presume that one will value happiness, and that they will choose the lifestyle that they think makes them happiest. However, people may hold different mental models about what makes them happy. If one holds a mental model that sees their career as the most important thing for their happiness, that sees healthy food as not worth the money, and leisure not worth the time, then this mindset will ultimately result in a lifestyle that promotes career over health and leisure. If, on the other hand, we have a mental model in which we believe health and leisure are as important to happiness as one's career, then we will lead a totally different and more health conscious lifestyle. This will ultimately lead to different patterns and events in our life - for example, getting less colds!

The Iceberg model can serve as a useful way of looking at the world. It reminds us to look beyond the superficial level of events to see the patterns, system structures, and mental models that have ultimately caused those events. This gives us a deeper level for addressing issues.

As said previously, reacting to events is like treating the symptoms of an illness without trying to treat its causes. Recognising patterns can help us anticipate events, but this in itself does not give us power to control those patterns or events. Changing the patterns requires us to work on the level of system structure, which in itself is something that almost always requires us to work on the level of mental models.

For a more enlightening example than catching a cold, consider the collapse of the Lehman Brothers bank in 2008. This event was part of a larger pattern whereby financial institutions who were heavily invested in the mortgage industry at the time suffered huge losses, having previously made huge profits. Systems thinkers would look deeper here and ask about the system structure that caused this pattern of growth and collapse. They would ask about the configuration of relationships and incentive structures between regulators, banks, homeowners and other actors, as well as the values held by actors in the system.

In order to prevent future severe recessions from occurring, it is clear that we would need to re-design the structure of the mortgage industry, and perhaps of many parts of the financial system. This in itself will require a significant shift in the mental models of not only policy makers and regulators, but also the everyday citizen. According to systems thinking, acting on this level of the system is what will produce the most successful and long term change.

As such we can see that it is our mental models that ultimately determine the way a system operates. The most effective way to change a system, then, is to change the mental models of the actors within it. This is really the ultimate goal of system dynamics. Rather than simply revealing a systems structure, a system dynamics analysis should be considered successful when it helps people adopt more accurate or useful mental models about the world.

This often requires us to first model the way that the actors in a system perceive that system. Once these mental models are revealed, we can ask ourselves whether our perceptions match the reality. If they do not, then we are offered an opportunity to change our perception (i.e. our mental model) so that it better reflects reality. This in turn leads to better decision making that is based on a more accurate picture of the world.

Many people would argue that a linear, reductionist, "too-narrow" view of our world still dominates society, whereby we do not account for the wider system within which many problems are embedded. In complex systems such as social systems, this often leads to policies and actions that worsen the very problems they are designed to resolve. System dynamics offers us a methodology for recognising the actual complexity of real world systems, rather than shying away from it. This can lead us to discover better and more long term solutions to problems, and to better communication of the logic of those solutions.

The Endogenous Point of View

We can now move to the second tenet of system dynamics, which is very much related to the tenet that structure drives behaviour. The endogenous point of view relates to the boundary that we set when analysing problems (whether through building models or another method of analysis). It asks us to include within our boundary of analysis all relevant variables that contribute to the problem at hand, rather than simply saying that the problem is occurring due to external (exogenous) factors that are outside of our control.

Some examples might make this clearer. In Industrial Dynamics, Forrester showed how General Electric's internal hiring policies were the cause of the three year cycle in employment levels, rather than something outside the business, such as fluctuating demand in the economy. In this sense, Forrester showed that the problem was being generated endogenously (within General Electric), rather than exogenously (outside General Electric). By creating a model which purposely set the demand in the economy as constant, Forrester was demonstrating how the problem (fluctuating need for staff) was not caused by the economy, but by the managers' hiring policies.

Finding that the source of a problem is within a system rather than external to it is good news. This is because if the source of the problem is within the system that we can control, then it is within our power to change it. In the case of the General Electric managers, they were able to identify how their own policies caused the three year cycle, and so they could alter their policies to avoid this. If the problem was thought to be external to the system (i.e. the cycles were caused by fluctuating demand in the economy), then General Electric would have had less agency to alter the employment cycles.

Implications of the Endogenous Point of View

Adopting the endogenous perspective has some important implications. Many of the big issues today can only be tackled when looked at from this perspective. Take climate change as an example. The exogenous perspective on this issue is that the Earth's rising temperatures are the result of natural cycles outside of human control (such as variations in solar radiation), meaning that little can be done to change this. The endogenous perspective, on the other hand, would recognise the role that humans have been shown to play in this issue (via the burning of fossil fuels, etc.) and would therefore advocate for certain policies to prevent further climate change, such as subsidies for renewable energy.

We also saw an example of the endogenous perspective in the Great Lakes case, which showed how the policies for constructing protective structures to counteract the adverse effects of flooding could actually lead to more housing being developed on that land, and therefore more vulnerability to such adverse effects. This adopts an endogenous perspective on a situation where at least one major factor, the water fluctuations of the lakes, was outside of the community's control.

An exogenous perspective in this scenario would see the flooding as something that simply happens, and that when it does, one must simply recover and rebuild. It does not recognise the part that human policies and actions play in the process.

A table that summarizes the differences in perspectives and corresponding policy implications can be seen below. This was taken from George Richardson's 2011 paper "Reflections on the foundations of system dynamics", which we recommend as a source of further reading on this topic.

Flood Damage	Perspective	Policy Implication
Exogenous view	Floods happen sometimes; the greater the flood, the worse the damage	When floods happen to occur, recover and rebuild
Endogenous view	Damage occurs when hazard meets vulnerability; vulnerability is a result of people's policies	Recognise human role in damage. Work with stakeholders to minimise vulnerabilities.

Summary of exogenous and endogenous agency perspectives on flood damage

Sections Summary

Having discussed some of the central tenets of system dynamics, we can now begin to consider some of its limitations. This is the focus of the next section.

Section 2.5: The Limitations of System Dynamics

In this section we will look at what it means to model a system, and this will lead us to some of the limitations of modelling in general. We will also reflect on the particular limitations of qualitative system dynamics modelling in comparison to quantitative system dynamics modelling.

Model Definition

System dynamics is a technique that uses models to increase understanding of a system. We have been using the word "model" quite often in this course so far, and so it may be worth taking a moment to define it. Our definition of a model is as follows:

A model is an abstraction (or a human construct or representation) of reality to achieve a certain objective.

Aside from mental models, there are many other types of models, used in our daily lives to make sense of the world. Some of us have used solar system models in elementary school to understand how planets align around the sun. Some have used maps to navigate through a new city. Others might have simply drawn human stick figures to communicate a message.

The Limitations of Models

It is important to understand that models are constructed to represent reality, and are not reality itself. As is often said, “the map is not the territory” (if it was, it would not be a map but just a big piece of paper that covers everything!). To be useful, models always include some sort of simplification of the real world. This brings us to the conclusion that:

No model is a perfect representation of a reality

In addition to being simplifications of reality, models are also the product of an observer with subjective understandings of the world. This leaves them open to biases. Nonetheless, although models never reflect reality 100%, good models can still be extremely useful in achieving the objective they were designed for. A map is not the territory, but a good map still helps us to arrive at our destination. Newton's theory of gravity is not completely accurate, but it still helps us to predict the motion of large bodies with incredible accuracy.

Similarly, CLDs cannot describe every cause and effect relationship in a system; if they did they would be too complex to understand, and would not achieve their objective of simplifying complexity and increasing understanding or communicating insights. However if the right level of abstraction is chosen, then the model can achieve its objective, whether that be communicative, analytical, or both. This leads us to the more nuanced conclusion that:

No model is a perfect representation of a reality, but good models can help us to better understand and better navigate that reality

The Limitations of CLDs

As mentioned in the Introduction, we will only be focusing on qualitative systems models (i.e. CLDs) in this course. We will not be focusing on quantitative (i.e. simulation) models. By simulation models, we mean models that assign values to variables in a model, and use mathematical equations to describe the relationships between related variables. Using specific software, one can produce graphs of how the different variables should co-develop over time, based on the assumptions in your model.

CLDs have advantages in that they are easier to construct and can be more easily understood in comparison to simulation models. In many cases, building a CLD is sufficient for the task at hand. However, in many cases it is often necessary to build simulation models to develop a sufficient understanding of a system. There are several reasons for this, and it is good to be aware of them. We will discuss each disadvantage in the pages that follow.

Understanding Dynamic Complexity

One disadvantage of CLDs when compared to simulation models is that CLDs are not as helpful in understanding how variables might interact overtime. When there are several variables in a model, our minds simply do not have the cognitive capacity to understand how those variables might co-develop over time, given the numerous relationships between them, many of which are non-linear. Simulation software does these calculations for us, in the same way that a calculator helps us do sums that are too challenging to do with our minds alone.

Nonetheless, we should always remember that simulation models are also limited - they do not simulate the future, as they rely entirely on the assumptions we put into the model.

Model Validation

Another advantage of simulation models over CLDs is that simulation models allow us to test our model against reality. If we fill our model with numeric data and then simulate it, we can see if the simulations of the model match the actual historic data. For example, if we build a simulation model of the population growth of deer on an island from 1990 up to 2030, we can see if the model's simulation of the deer population matches the actual historic data (i.e. the data from 1990 until present), which gives us a good indication about whether or not the model can be used as a good tool for simulating future scenarios of deer population development on the island. Given that CLDs do not use numeric data, this means they cannot be tested in this way, which means there is less opportunity to test their validity.

End of Section 2.5: The Limitations of System Dynamics

Despite these limitations, CLDs can still be very useful tools in many contexts, by helping individuals and groups to better analyse problems and communicate their understandings and assumptions. They also serve as a good first step towards building a simulation model, and can help one to determine whether or not it is necessary to build a simulation model at all.

Review of Module Two

That brings us to the end of Module Two. To recap, we began the chapter with a series of complex questions, and asked you to spend some time thinking about some of these and writing down your answers. We then explained how systems thinkers would say that these answers are based on your mental model about those topics/problems. From there we discussed what the term mental models means, and talked about how it can be difficult to become aware of or change our own mental models about complex topics. We also discussed how communicating our mental model to others can be very challenging.

System dynamics was then proposed as a tool to help us formally model ours and others' mental models. It was said that system dynamics models can help us to both develop and communicate our mental models about complex systems. In other words, they help us to develop and communicate our understanding of how the world works.

Following a brief overview of the origin and history of system dynamics, we touched upon two of its central tenets: that structure drives behaviour, and that it is better to adopt the endogenous point of view when addressing problems. After reviewing the benefits of system dynamics modelling, we noted some of its limitations, and particularly the limitations of qualitative as opposed to quantitative modelling.

We hope that by now you are excited to learn about system dynamics. In the next modules we will begin demonstrating what CLDs look like, how to read them, and how to build them. More importantly, we will demonstrate how building these models can help us to better understand complex problems, and better identify the most effective and long-term solutions to our and our society's challenges.

System dynamics is essentially just a technique. However, the use of this technique can radically transform the way in which one approaches problem, and even the way in which one understands the world. We hope that this course will not only offer you a new technique, but also offer you an opportunity to look at the world in a new and hopefully clearer light.

Module Three
The Basics of Systems Modelling

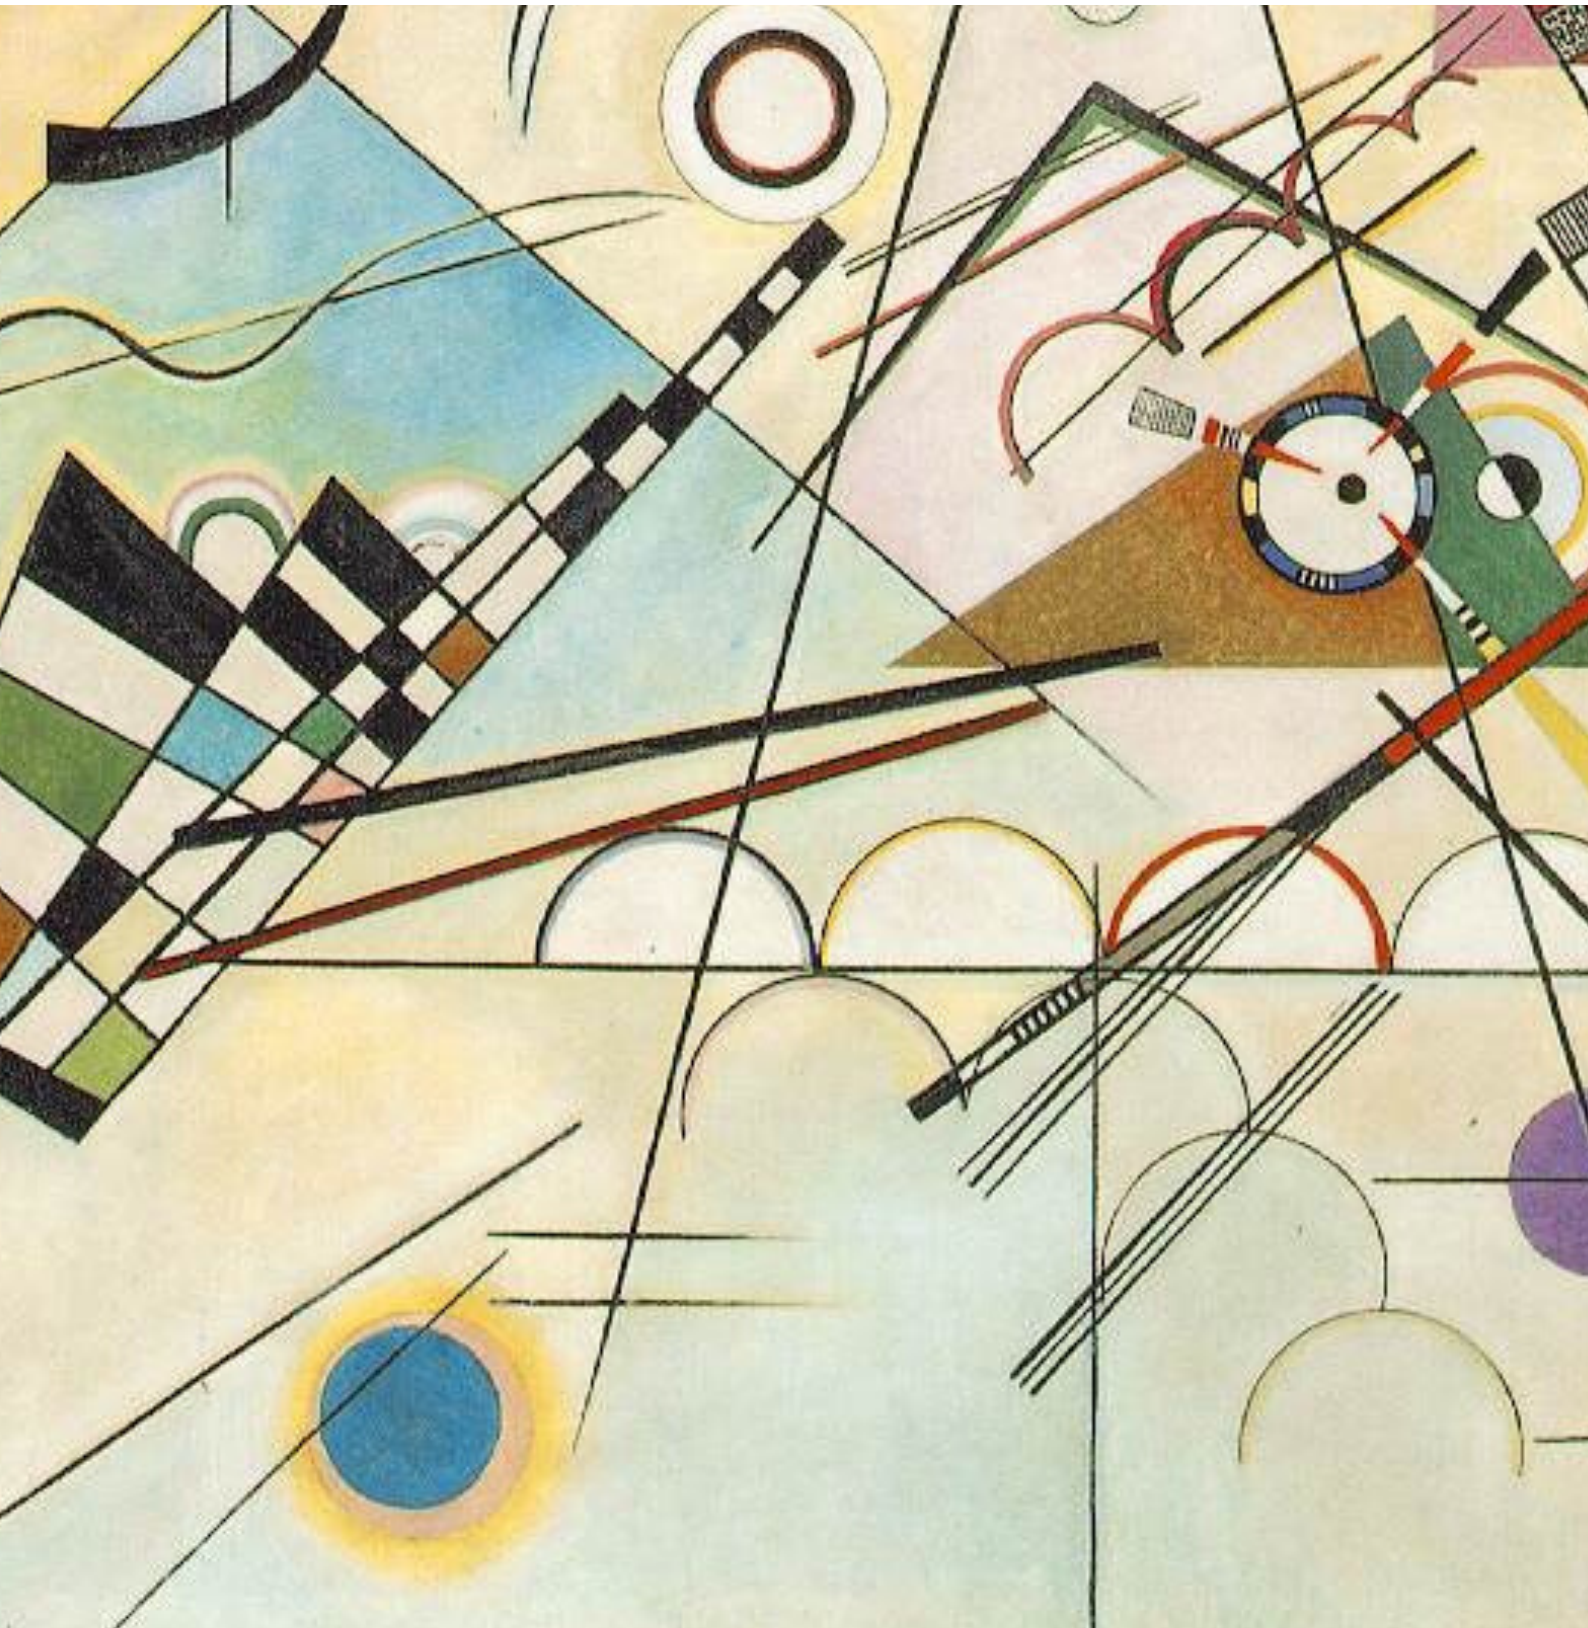

WASSILY KANDINSKY - CIRCLES IN A CIRCLE, 1923

Module Three Overview

In the previous module, you have learned that we build CLDs to help us understand and find solutions for issues in complex systems, as well as to communicate our assumptions more explicitly and show the logic behind our thinking and/or any solutions we propose.

In this module, we will focus on the basic building blocks of CLDs. Explaining how they use words and arrows to display the different cause and effect relationships and feedback loops that are assumed to give rise to a certain problem or pattern of behaviour.

Module Outline

The module is divided into four sections, as described below:

1. Casual Relationships in CLDs

2. Feedback Loops

3 Behaviour Over Time Graphs

4.Systems Archetypes

Module Goals

By the end of this module you should be able to:

Read causal relationships in CLDs.

Understand and identify the different kinds of feedback loops (i.e., reinforcing and balancing) in a model

Read behaviour over time graphs

Identify and describe systems archetypes and leverage points

Section 3.1: Causal Relationships in CLDs

If you consider casual loop diagramming to be a language, then you can think of the next section as teaching the basic grammar of this language. And while grammar is not always everyone's cup of tea, it is a necessary step on the way to speaking and understanding a language fluently. So please bear with us if the next few pages aren't so captivating; it will be worth it eventually!

Two Types of Causal Relationships

A cause and effect relationship in a CLD is represented by an arrow between two variables. By “variable,” we mean literally anything that can change over time. The arrow goes from the cause variable (the one causing the change) to the effect variable (the one being affected by the change).

There are two types of cause and effect relationships in CLDs: positive and negative. A positive relationship is depicted by the '+' symbol at the arrow's head, whereas a negative relationship is depicted by the '-' symbol at the arrow's head. In this course we will use also use blue arrows for positive relationships and red for negative ones, although this is done by all system dynamicists.

Note: Remember, to zoom in and out of model pages, use the buttons in the top right corner of the page. The symbol with two arrows will fit the model zoom to capture the whole model.

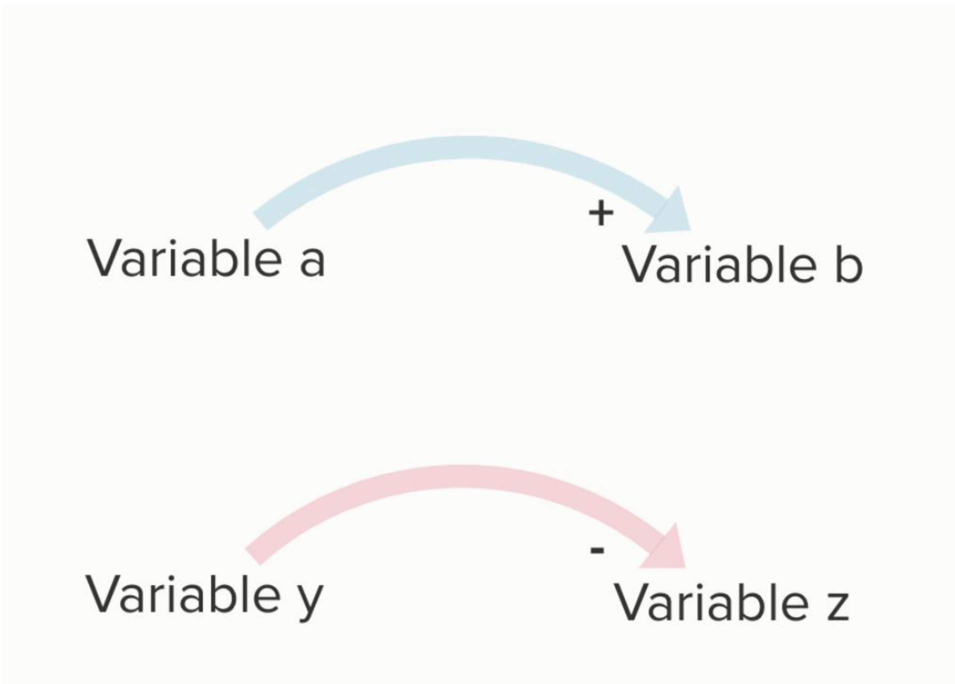

A positive relationship means that the cause and effect variable are moving in the same direction, such that an increase in the cause variable results in an increase in the effect variable, and a decrease in the cause variable results in a decrease in the effect variable, all else equal.

Note that the term positive does not mean that the relationship is necessarily a desirable one. Instead the term positive here comes from its mathematical use, and only means that the two variables move in the same direction.

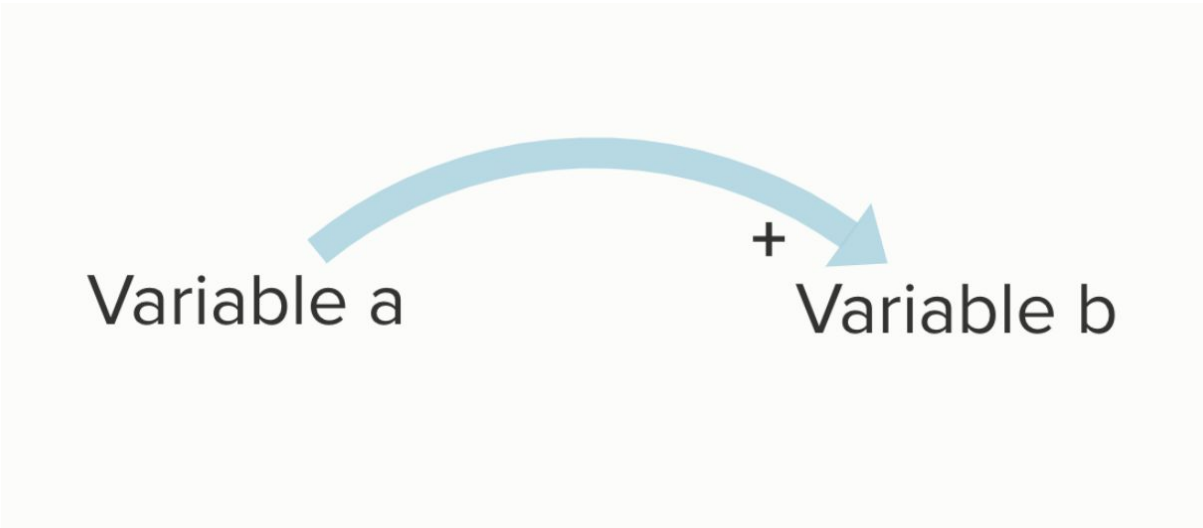

Conversely, in a negative relationship, the effect variable moves in the opposite direction to the cause variable. So in this case an increase in the cause results in a decrease in the effect, and a decrease in the cause results in an increase in the effect, all else equal.

Again, the term negative does not mean undesirable but rather refers only to the fact that the two variables move in the opposite direction.

Some examples will help make things clearer:

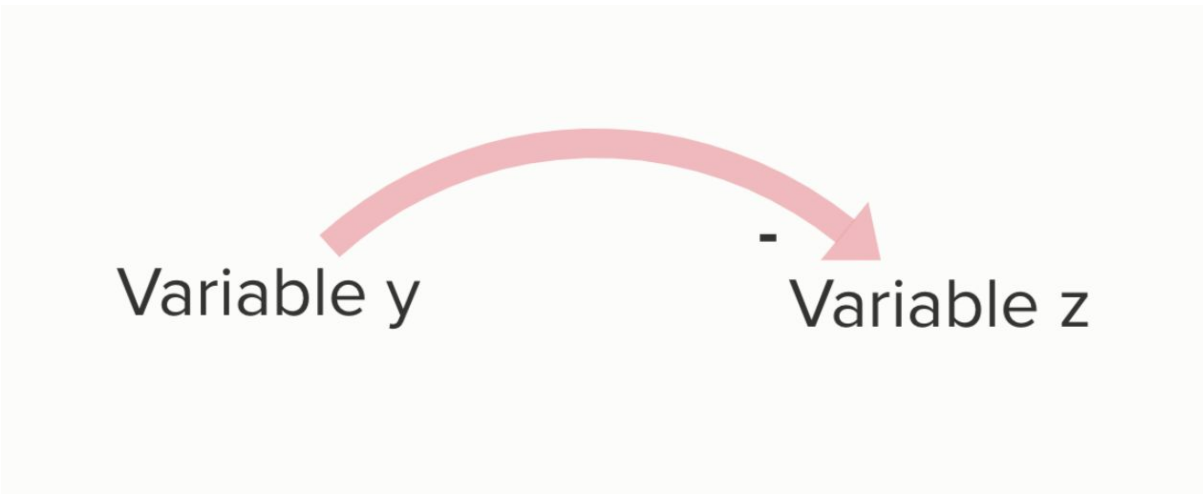

Example of a Positive Relationship

You may remember the two mental models we expressed in Module Two relating to the question of how lowering working hours might affect a country's material wealth. In the first mental model, the opinion was expressed that lowering working hours will result in less material wealth because there will be less goods produced and less services provided.

The causal relationship expressed in words here can be formally modelled as seen below. This causal relationship is saying that if the Hours worked per week were to be decreased, then there would be less Goods produced/services provided per week. It is also saying that if the Hours worked per week were to be increased, then the Goods produced/services provided per week would also increase, all else equal.

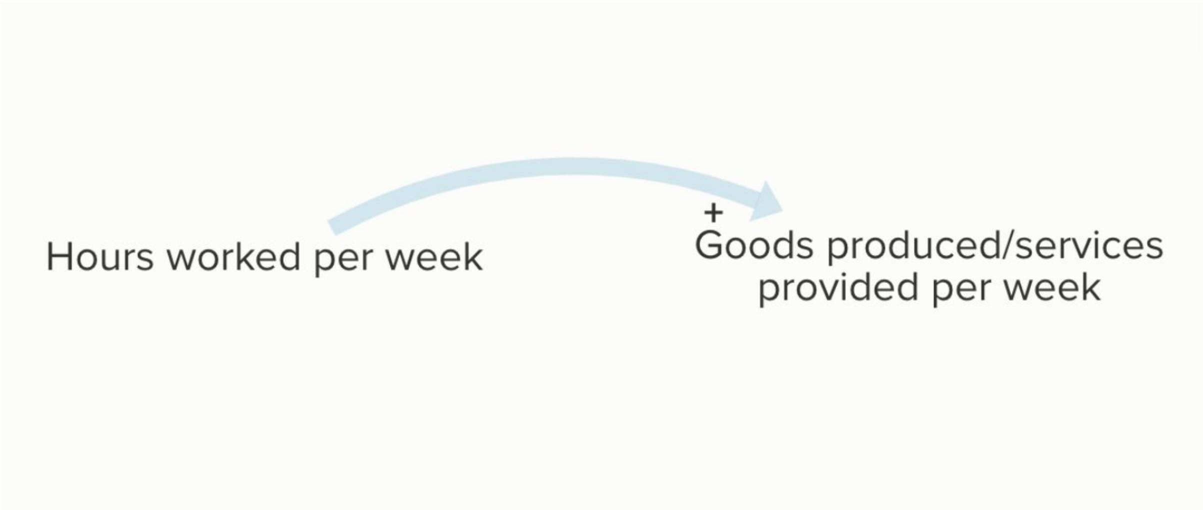

Another example of a positive relationship is Births to Population. This is a positive relationship because an increase in Births will mean that the Population will be larger than it would be if the Births hadn't increased. Additionally, a decrease in Births will mean that the Population will be smaller than what it would otherwise have been.

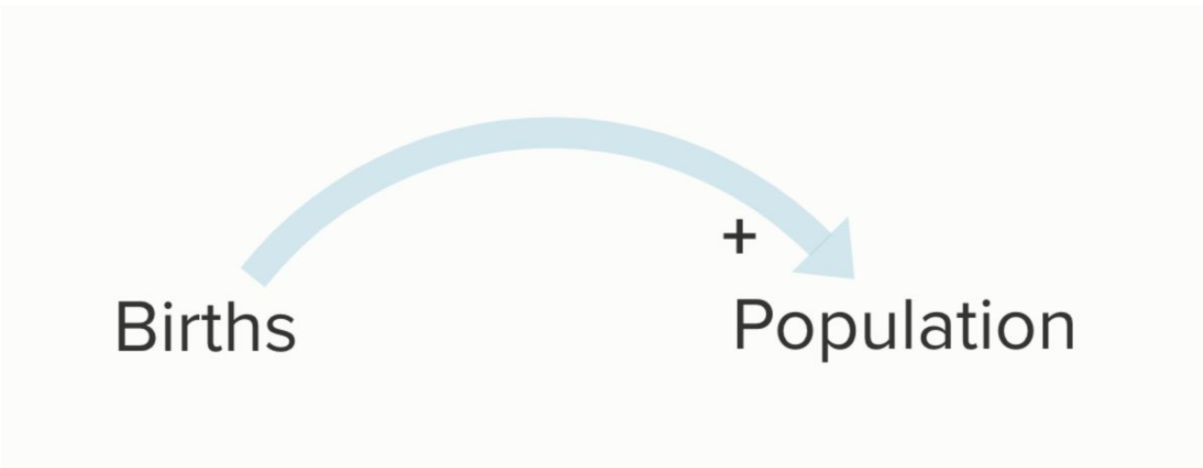

... what it would otherwise have been

On the previous page we used the phrase “...what it would otherwise have been". This phrase is applicable to all causal relationships in CLDs. We say this because a decrease in the cause variable does not necessarily mean a net decrease in the effect variable (in a positive relationship). It just means that the effect variable will be less than it would otherwise have been (i.e. if the cause variable hadn't decreased). Additionally, in a positive relationship an increase in the cause variable does not necessarily mean that the effect variable will increase, it just means that the effect variable will be more than it would otherwise have been (i.e. if the cause variable hadn't increased).

For example, if the births in an area decrease from 10,000 a year to 7,000, the native population will still continue to grow as long as births remain greater than deaths in the area. However the decrease in births means that the population will grow at a slower rate. As such, the decrease in births (the decrease in the cause variable) means that the population (the effect variable) will be lower than it would otherwise have been.

All Else Equal

You may have also heard us use the term all else equal when describing the meaning of causal relationships. This is because causal relationships in CLDs are always considered in isolation to the rest of the system. When we determine each causal relationship, we always assume that all other variables in the system will remain equal. You may hear some systems thinkers use the term "ceteris paribus", which is Latin for "all other things being equal".

So based on the model below, we can say that more Access to healthcare will lead to a greater Life expectancy, ceteris paribus.

Of course, Access to healthcare is not the only variable that affects Life expectancy. For example, one's Quality of diet will also affect their Life expectancy. This can be modelled as seen below:

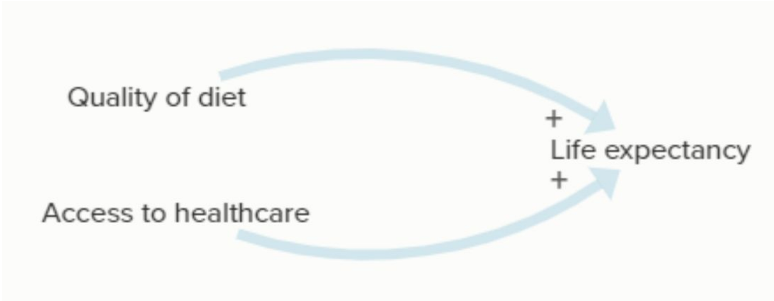

As such, when reading or developing causal connections in CLDs, we always determine/read the type of relationship in isolation to the other variables (i.e., assuming that they remain equal). So the above model is saying that you an increase in Access to healthcare will increase your Life expectancy, without regard to what is happening to your Quality of diet, which is a separate variable and so modelled separately.

Exercise 3A: Positive Relationships

Can you think of a positive relationship? It can be related to your personal life, to your own studies or profession, or to something that occurs in nature.

The important thing to keep in mind is that an increase in the cause variable will result in the effect variable being more than what it otherwise would have been, and vice versa.

Some examples of positive relationships include those summarised in the table below.

Cause variable	Effect variable
Access to electricity	Hours of the day you can read/study
Consumption	Resource use
Productivity	Finished tasks

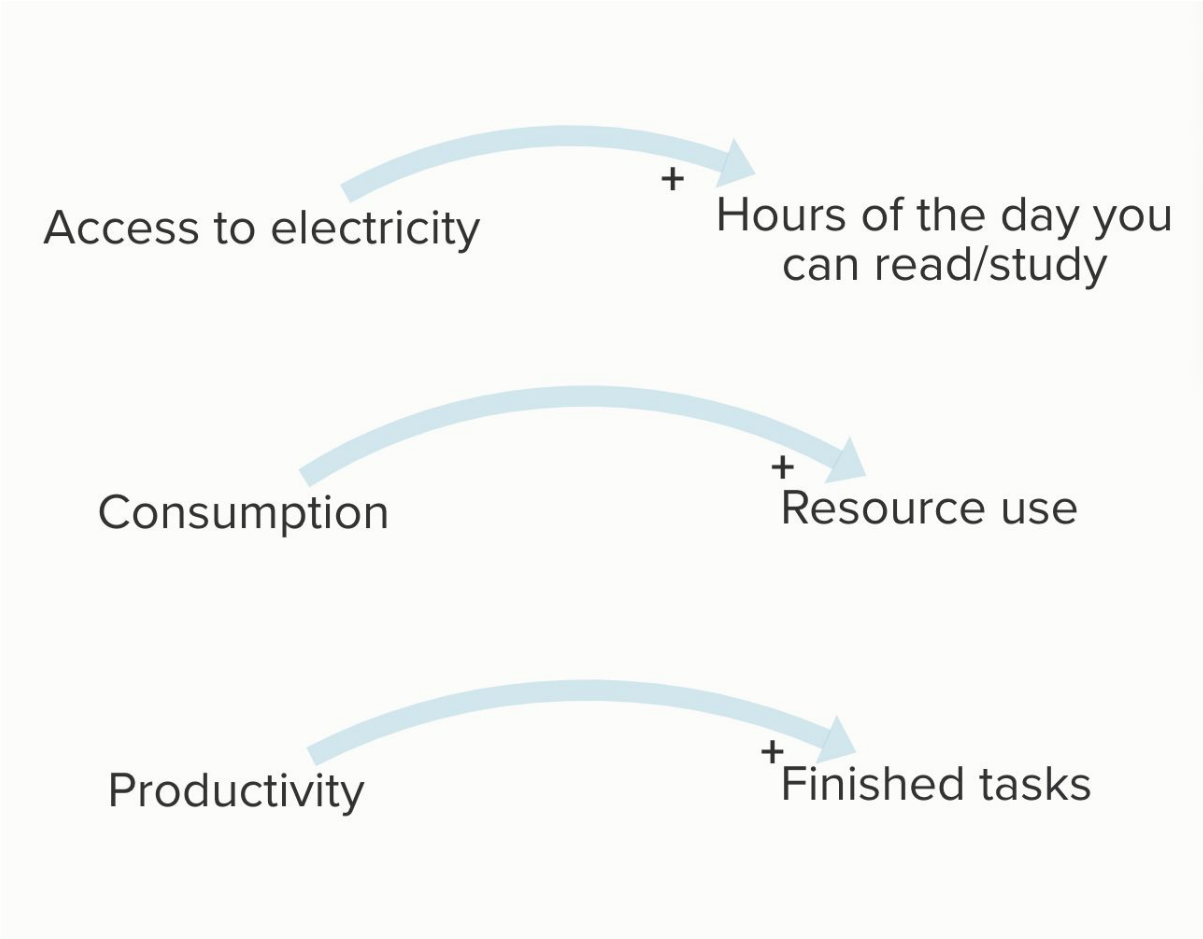

Example of a Negative Relationship

Below, you can see an example of a negative relationship. Here we are saying that if the Price of grapes increases, then Demand for grapes will decrease (due to some assumed price sensitivity of people who buy grapes). We are also saying that if the Price of grapes decreases, then the Demand for grapes will increase.

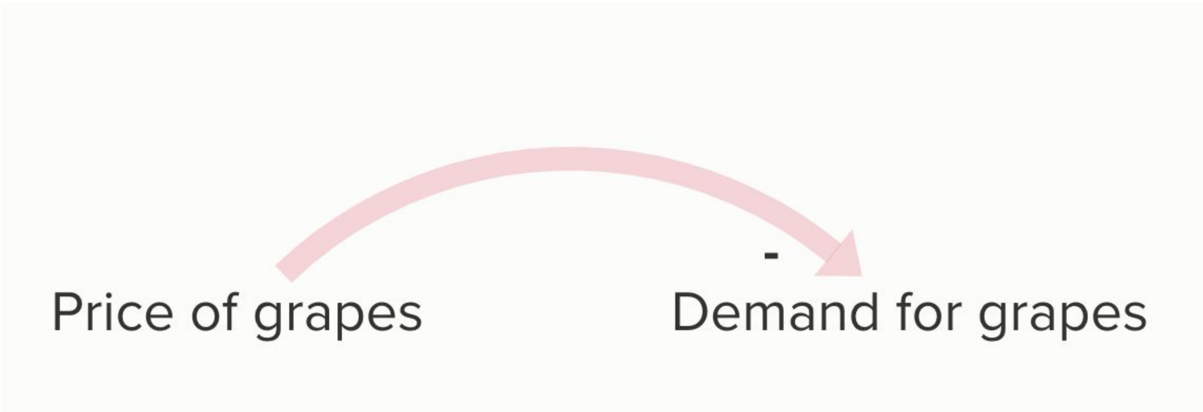

Another example of a negative relationship is that of Deaths to Population. An increase in the number of Deaths means that the Population will be smaller than what it otherwise would have been (i.e., if the Deaths hadn't increased).

If Deaths decrease, then the Population will be larger than what it otherwise would have been.

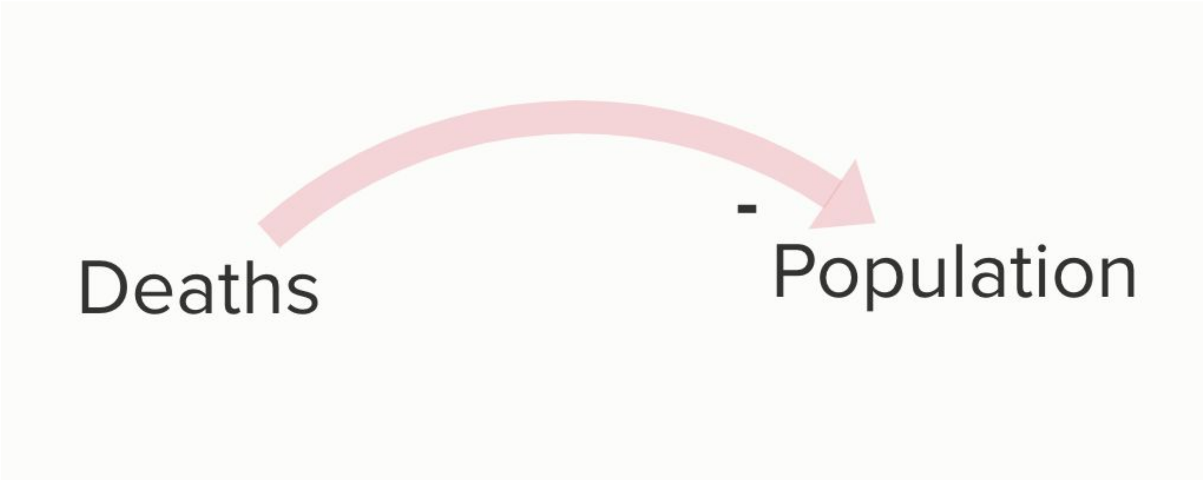

Exercise 3B: Negative relationships

Can you think of a negative relationship?

It can be related to something that occurs in nature, to your personal life or to your own studies or profession.

The important thing to keep in mind is that an increase in the cause variable will result in the effect variable being less than what it otherwise would have been, and vice versa.

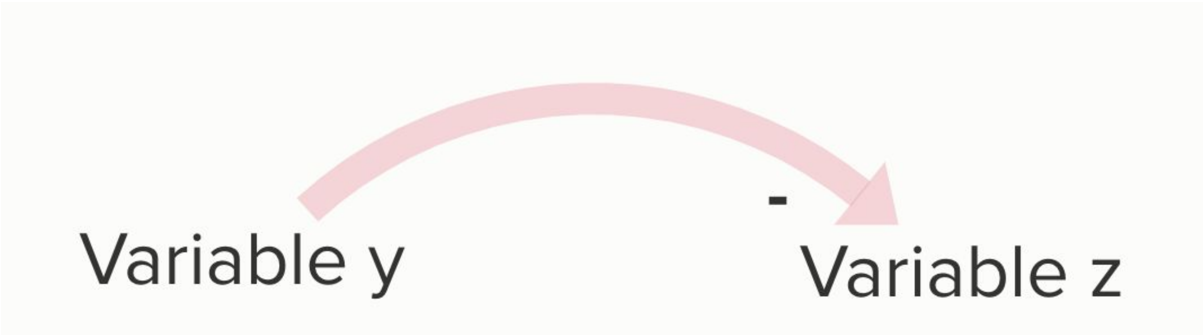

Answers to 3B

Some examples of negative relationships include those summarised in the table below.

Cause variable	Effect variable
Gender equality	Income gap
Environmental awareness	Meat consumption
Traffic speed cameras	Speeding

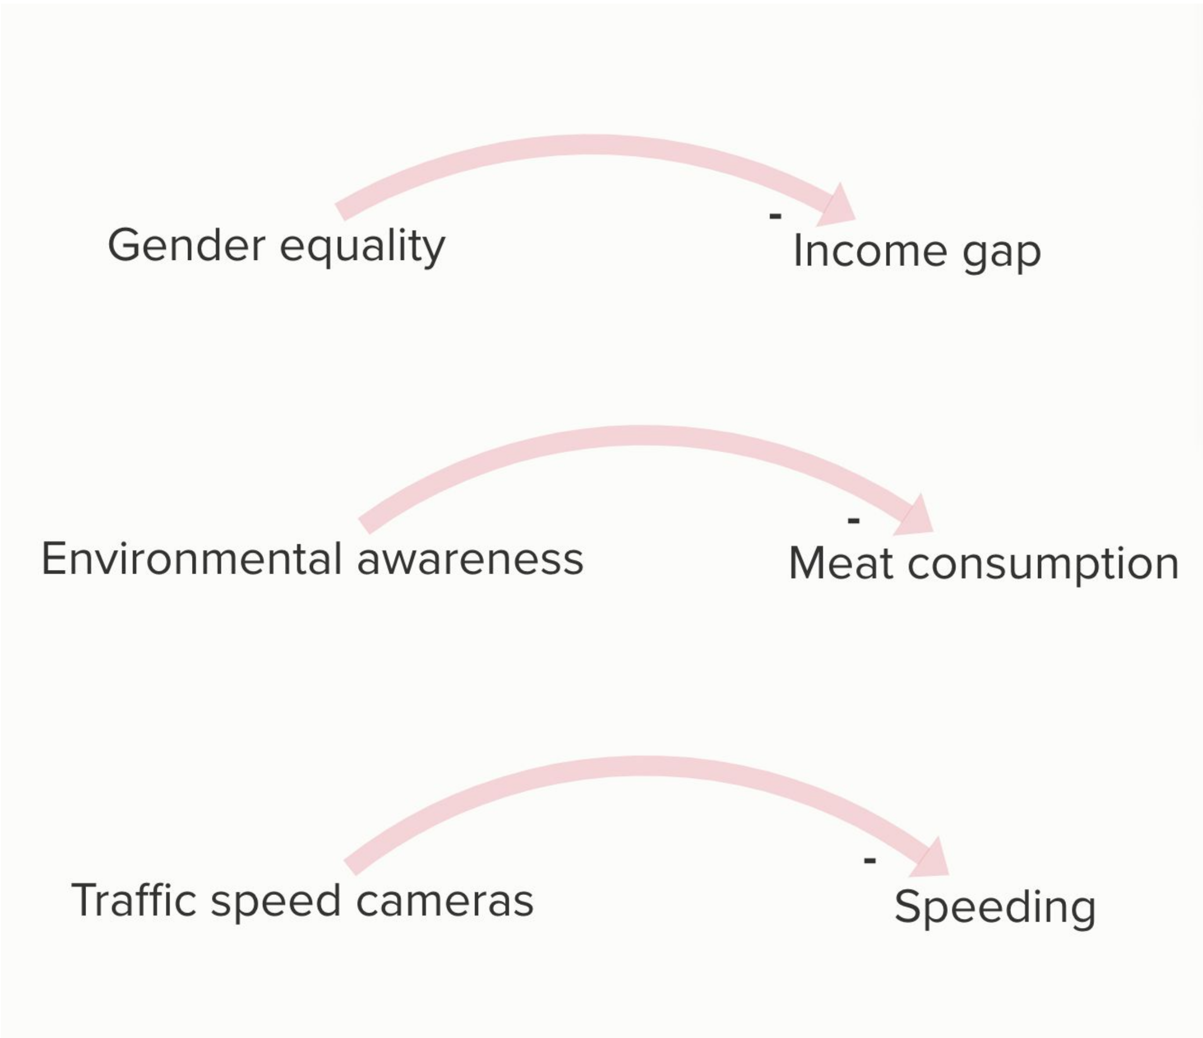

Causality not Correlation

Another important thing to note is that in CLDs, we portray relationships that we assume are causal, rather than those that are only correlated. An example might help us to illustrate the difference between correlation and causality: a famous study found that there is a strong correlation between ice cream sales and the murder rate on a given day. However there is no sensible way that we can imagine ice cream sales influencing the murder rate, or the murder rate influencing the sales of ice cream! A more sensible conclusion is that something else is at play.

Indeed, ice cream sales are likely to increase on a hot day. At the same time, studies have suggested that hot weather increases the murder rate on any given day, for a number of reasons. For one, more people are out of their homes, which increases the chance of conflict (for more on this theory, see this article). As such, it seems that temperature is the real causal variable that affects both ice cream sales and the murder rate.

Incorrect model:

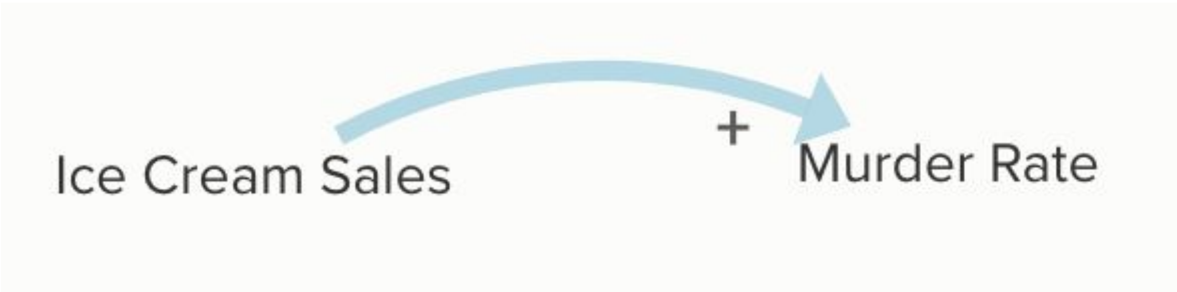

Correct model:

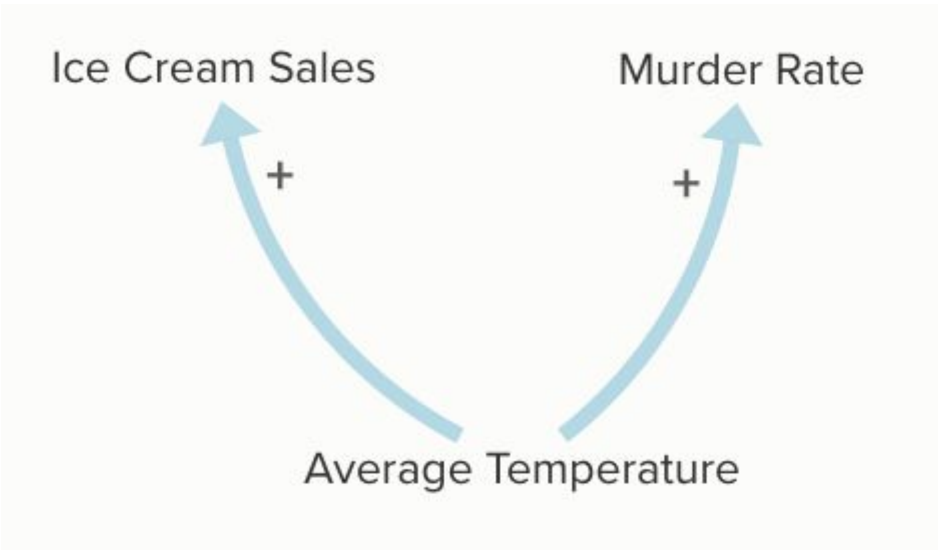

The important insight from this is that when creating CLDs, it is important that we only portray what we assume are causal relationships, rather than relationships of correlation.

Exercise 3C: Causal Relationships

Bellow you can see three variables (Interest payments, Withdrawals and Money in bank account).

Can you draw the arrows between the variables to depict the correct causal influences? And can you determine whether each relationship is positive or negative?

The answers are on the next page.

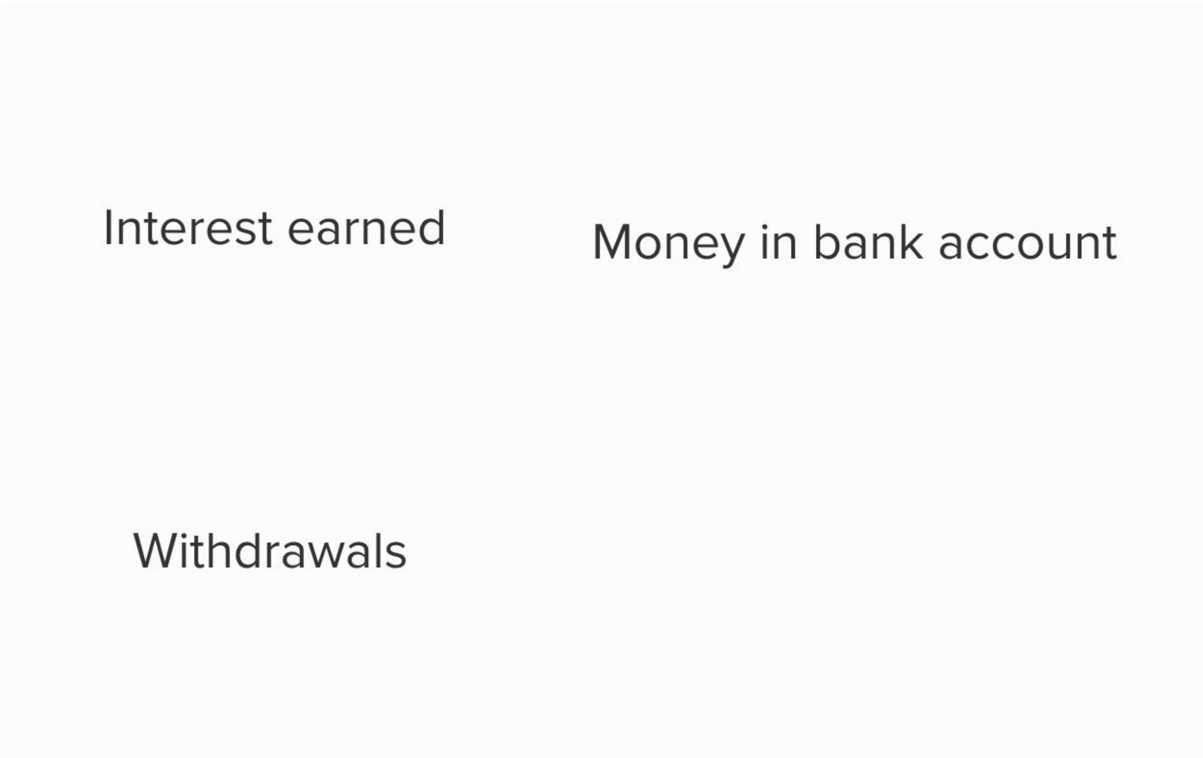

Answer to 3C

An increase in Interest earned (i.e. interest that you receive from the bank) leads to the Money in bank account increasing above what it would otherwise have been. Similarly, a decrease in Interest earned lead to less Money in bank account, compared to what would otherwise be the case. As such, Interest earned to Money in bank account is a positive relationship.

An increase in Withdrawals leads to less Money in bank account, and a decrease in Withdrawals leads to more Money in bank account, (compared to a case where Withdrawals hadn't decreased). As such, Withdrawals to Money in bank account has a negative relationship.

You may also have drawn a positive relationship from Money in bank account to Withdrawls, under the assumption that the owner of the bank account will take out more money when it is available. We will discuss such relationships later in the module!

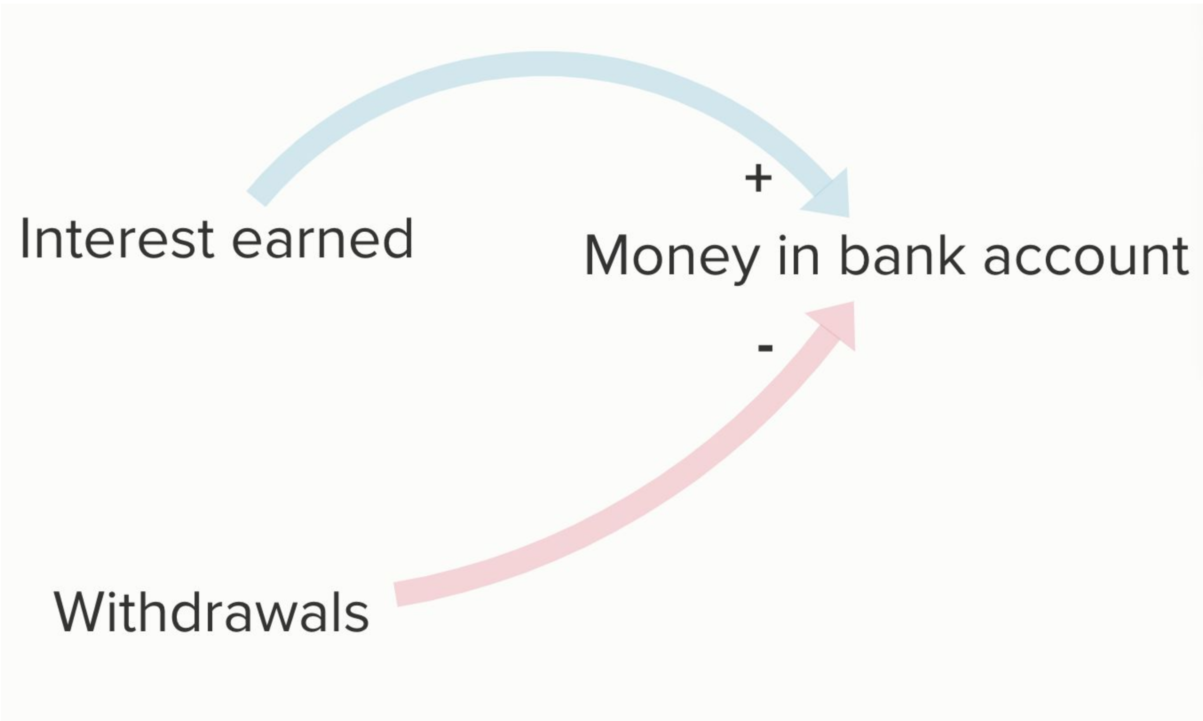

A Note on Naming Variables

When you start building your own CLDs, it is important to remember some best practices that will make your CLDs as useful as possible. Two of the best practices in naming your variables are:

Variables should be nouns or noun-phrases rather than verbs.
Choose variable names whose normal sense of direction is positive.
Let's explain why these are considered best practice.

1. Variables should be nouns or noun-phrases rather than verbs.

This is because the actions (verbs) are captured by the causal links connecting the variables. So for example, it is better that we have variables such as Costs and Price rather than variables with names like Costs rise and Price rises. Adding the verb rises to the diagram presumes that costs will only rise, which biases our discussion towards one pattern of behaviour (increasing costs and prices).

Not best practice:

Best practice:

2. Choose variable names whose normal sense of direction is positive

So for example, it is better to have a variable called Happiness than to have a variable called Unhappiness. This is because it is easier to think of an increase/decrease in Happiness rather than an increase/decrease in Unhappiness. And it is generally better to have a variable called Profits rather than Losses, because it seems clearer to talk about falling profits rather than falling losses.

Not best practice:

Best practice:

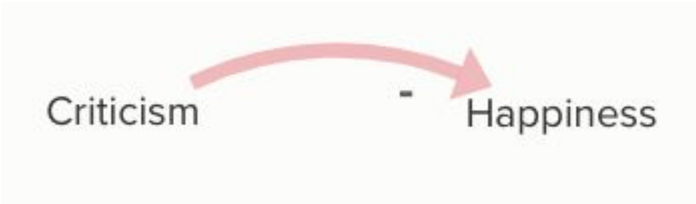

There may be some exceptions to this practice, but as one of the most prominent system dynamicists, John Sterman (of MIT), once said: "decreasing noncompliance with this principle will diminish your audience's incomprehension"!

Exercise 3D: Errors in Variable Names

Try to spot the errors in the variable names from the examples below:

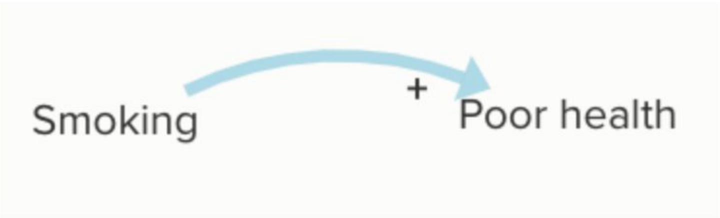

Answers to 3D

Not best practice:

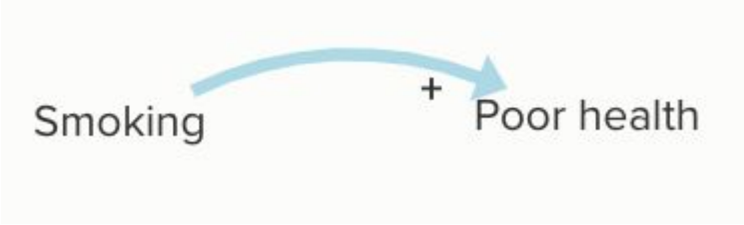

Best practice:

Not best practice:

Best practice:

End of Section 3.1: Causal Relationships in CLDs

That concludes the first section of this module. We hope that by now you:

Understand the meaning of positive and negative causal relationships
Understand that each causal relationship in a CLD is determined in isolation to the rest of the system (i.e., all else is considered equal)
See that relationships in CLDs are meant to signify causality between two variables, rather than correlation.
Are aware of best practices for naming variables
In the next section, we will begin looking at feedback loops.

Section 3.2: Feedback Loops

Now that we've gotten through the basic grammar of how to read and write causal connections in CLDs, we can finally move on to demonstrating how CLDs can help us to analyse a problem—in other words, to speak the language!

One of the most valuable things about constructing CLDs is that they help us to recognise feedback loops in the system at hand.

This section will describe what feedback loops are, and why it is so important to recognise them in a system.

Linear Chains of Cause and Effect

Up until now we have only modelled one-directional (also referred to as unidirectional) relationships between a cause variable and an effect variable. Don't worry, things will get more interesting!

In systems thinking, we of course recognize that effects are often causes to other variables. For example, more Organic farms leads to less use of Chemical pesticides, which positively contributes to Biodiversity, which in turn leads to higher Soil quality. Modelling these chains of causes and effects in CLDs can help us to recognize how certain variables are interrelated even though the relationship—or at least the mechanism by which that relationship exists— might not be obvious at first.

This is one of the benefits of causal loop diagramming. It aids our thinking, by having us map out causal relations that may otherwise be difficult to recognise, or to explain to others.

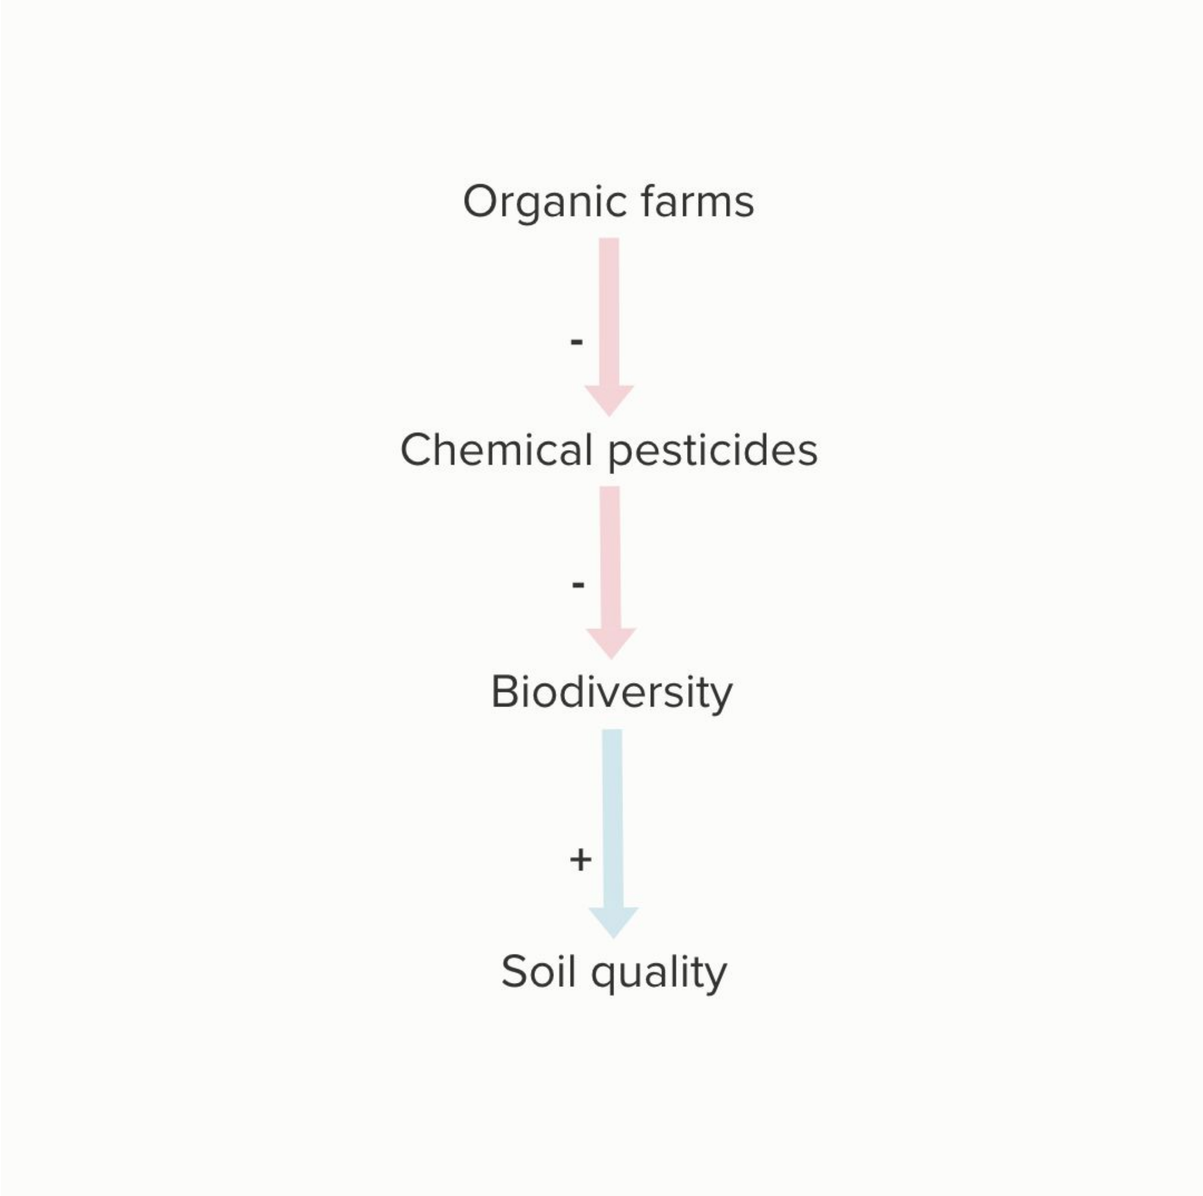

Circular Chains of Cause and Effect

When we follow a chain of cause and effect, we might often end up back where we started! Our last variable feeds back to the original cause. This means that we have discovered a feedback loop in the system.

A feedback loop is a circular chain of cause and effect. So when A affects B, B affects C, and C affects A again, we say that a feedback loop exists between these variables.

Just as there are only two types of causal relationships (positive or negative), there are also only two types of feedback loops: reinforcing or balancing.

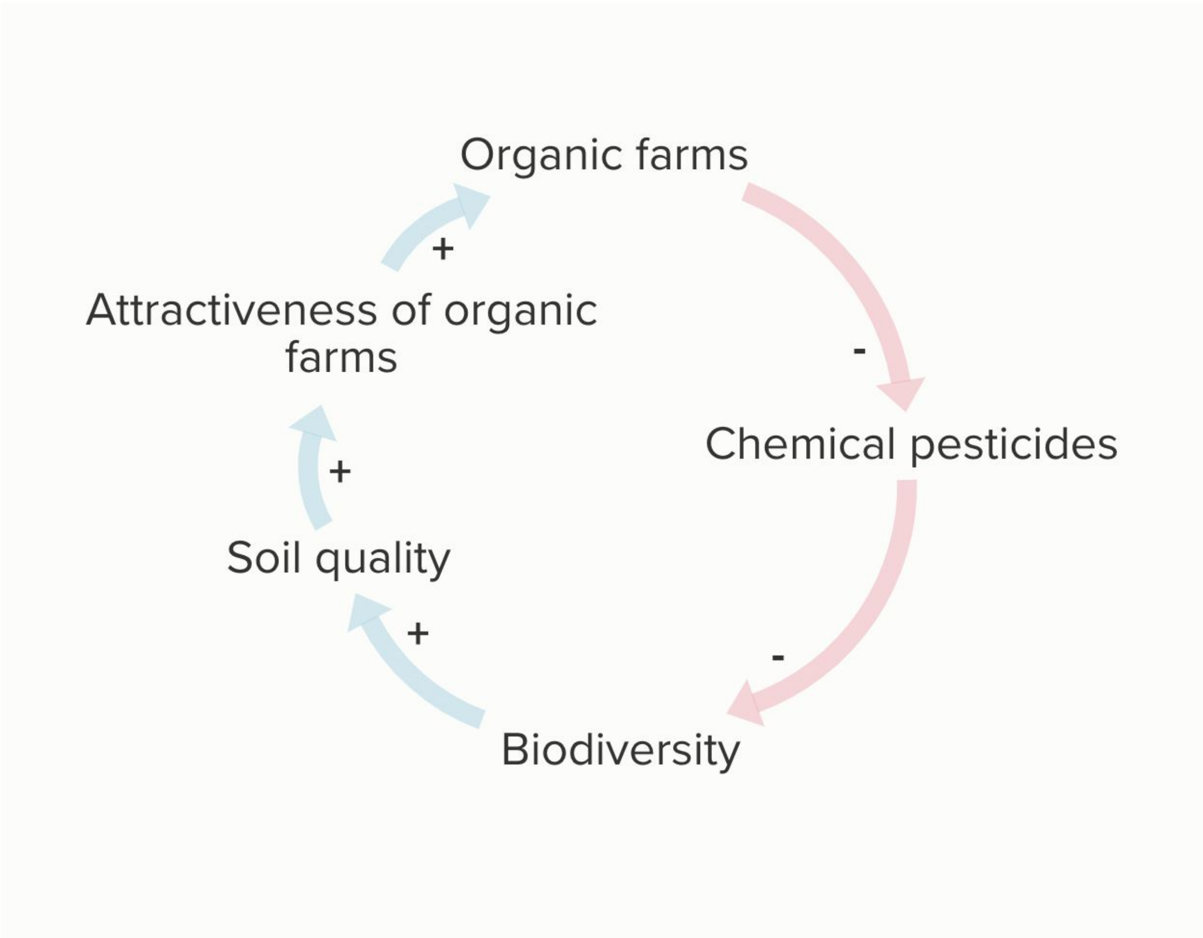

Reinforcing Feedback Loops (R)

An example of a reinforcing feedback loop is the one of Births and Population.

The positive relationship from Births to Population also works the other way around - the higher the Population, the more Births there will be, all else equal.

Therefore, an increase in Population will result in an increase in Births, and so a further increase in Population. The initial increase in Population is reinforced by this feedback loop structure in the system.

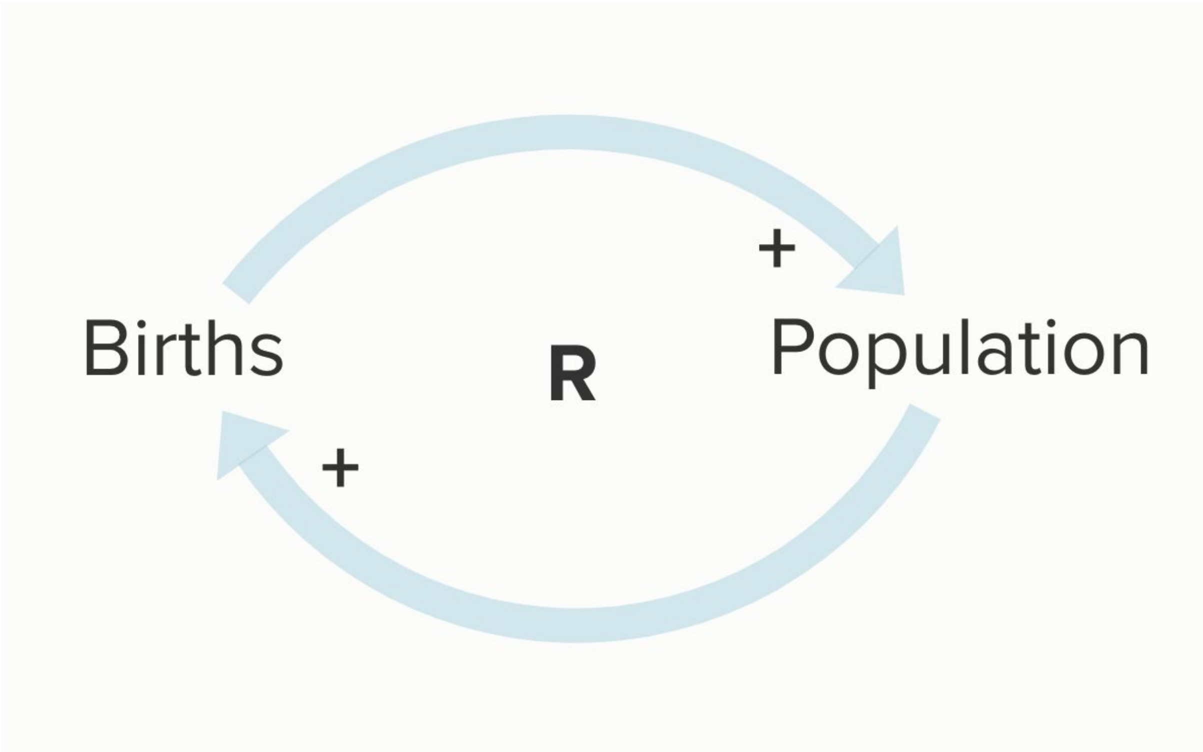

Reinforcing Feedback Loops (R)

Another example of a reinforcing feedback loop is that of the Ice Albedo effect.

An increase in Greenhouse gas (GHG) emissions per year leads to more GHGs accumulating in the atmosphere and more Heat trapped. Because the atmosphere gets warmer, more Ice melting will occur. More Ice melting leads to more open water which increases Sunlight absorption. More Sunlight absorption, in turn, leads to even more Heat trapped.

As you can see, reinforcing feedback loops are depicted with an R in the middle of the loop. If it's useful, we can sometimes give our feedback loops names, as we've done here.

Tip: to see all the variables that are part of a feedback loop, simply hover your cursor over the feedback loop name, and the relevant variables will be shown.

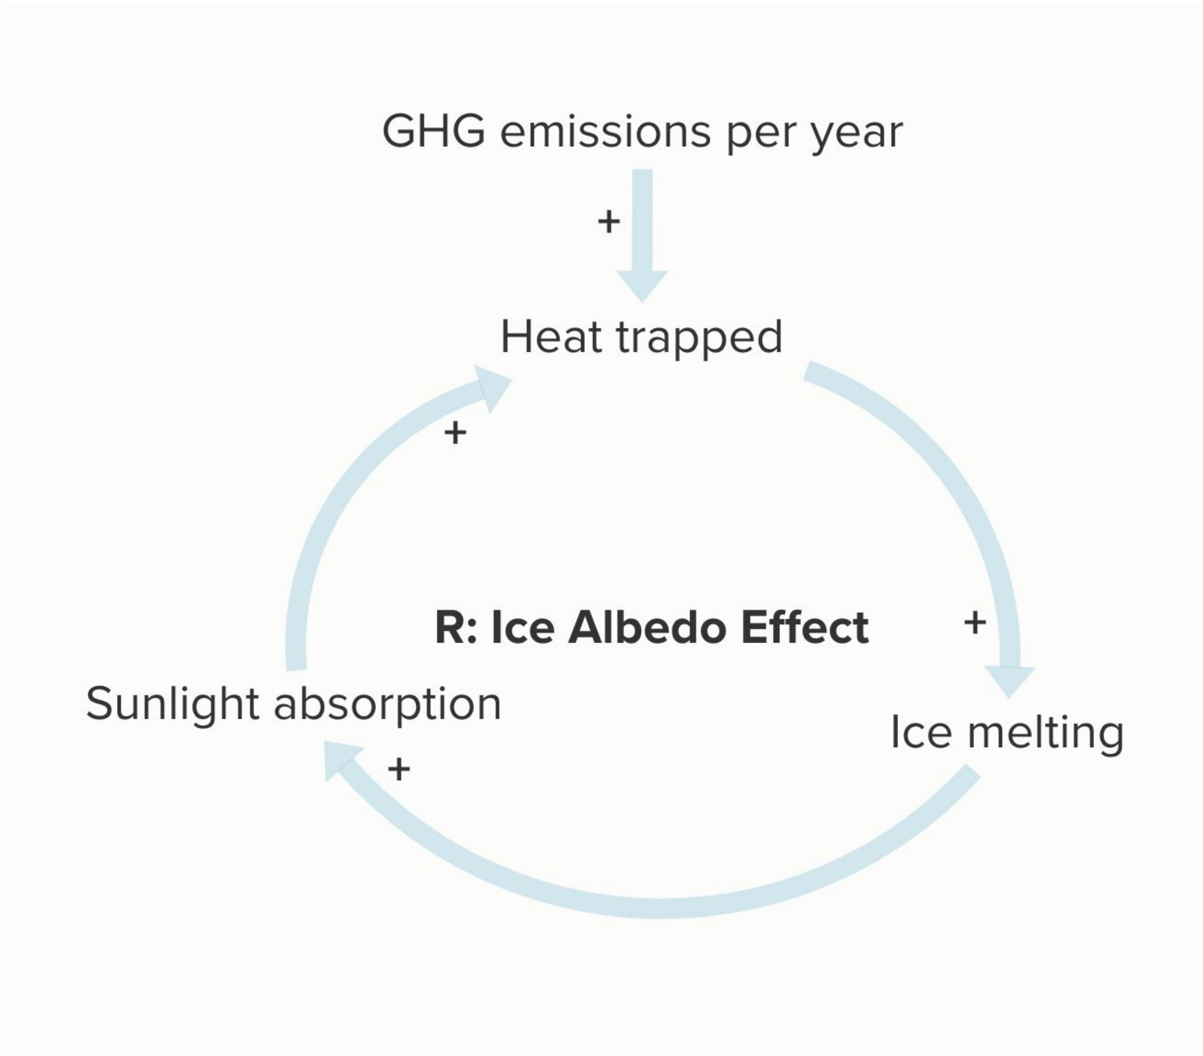

Reinforcing Feedback Loops (R)

To summarize, in a reinforcing feedback loop, a change of direction in one of the variables leads to even more change in that direction. As such, we can say that reinforcing loops always cause a greater amount of change in a system—they reinforce change, so to speak. This is why it is so important that we are always aware of reinforcing feedback loops in any system. The construction of a CLD is an effective way of identifying feedback loops in a system.

The change brought about by reinforcing feedback loops can sometimes be desirable, and sometimes undesirable. When the change is undesired, we say that there is a “vicious loop" (or vicious cycle) at play. When the change is in a desired direction, we say that there is a “virtuous loop” (or virtuous cycle) at play.

Exercise 3E: Reinforcing Feedback

Let's try and see if we can spot some feedback loops. Can you draw the correct causal connections and make a reinforcing loop with the variables shown below?

The answer is on the next slide.

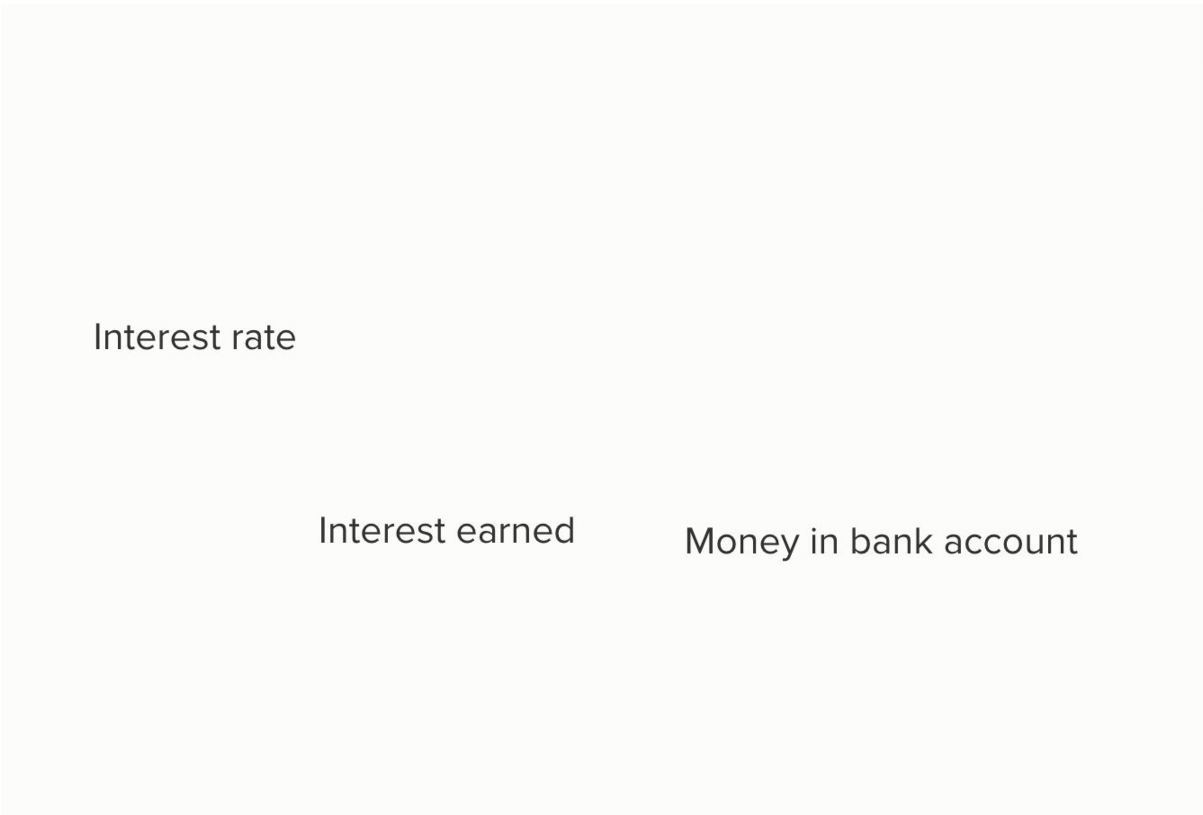

Answer to 3E

The Money in bank account determines the amount of Interest earned: the more Money in bank account, the higher the Interest earned. Higher Interest earned means more Money in bank account. As such there is a reinforcing loop between the two variables.

Conversely, if the Money in bank account reduces, then so do the Interest payments, which means even less Money in bank account, compared to what would otherwise have been.

Note that the Interest rate is considered constant in this example and that there are no withdrawals. Remember the ceteris paribus concept?

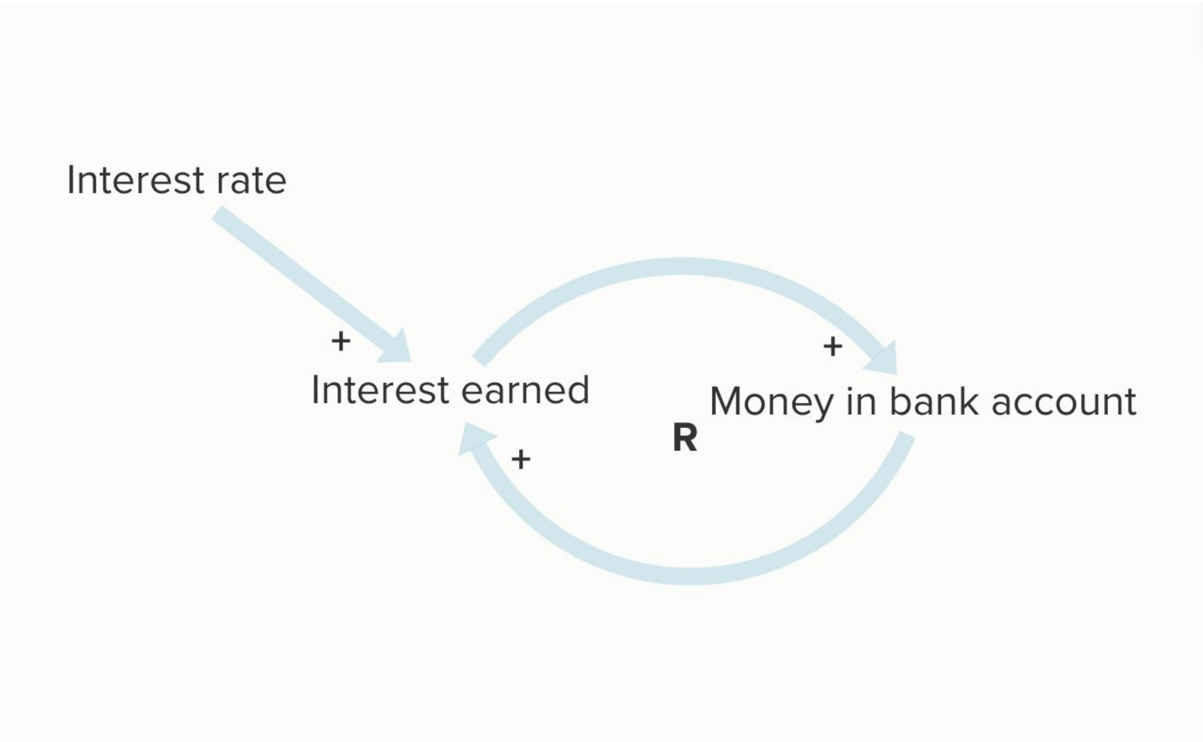

Balancing Feedback Loops (B)

The other type of feedback loop that exists is balancing loops. Whilst reinforcing loops generally cause greater change within the system, balancing loops generally do the opposite - they prevent change and create greater stability in the system, whether that be desired or not.

An example of a balancing feedback loop is that of Deaths and Population. When a Population increases, the total number of Deaths also increases, all else equal. With an increase in Deaths, the Population in turn decreases. An increase in Population is thus balanced by an increase in Deaths. Conversely, a decrease in Population will be partially compensated for by a decrease in Deaths.

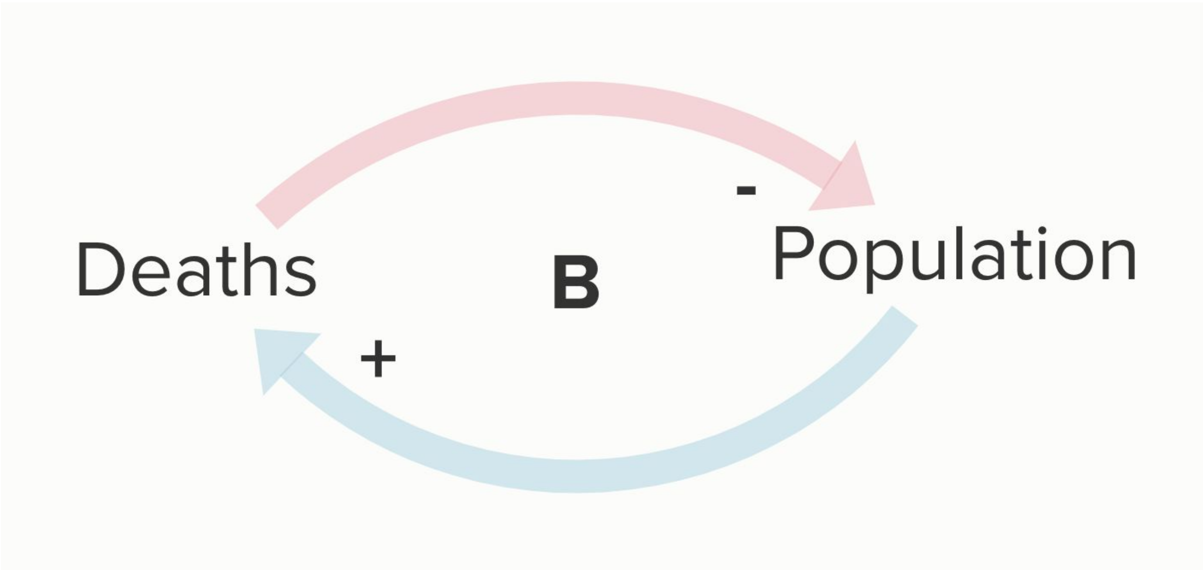

Balancing Feedback Loops and Goals

Whenever there is a Desired state in a system, then there will almost always be a balancing loop that describes how that goal is achieved.

We say that as the Gap between the Desired state and the Current state widens, there will be more corrective Action. This brings the Current state closer to the Desired state, hopefully until the Gap between the two is closed.

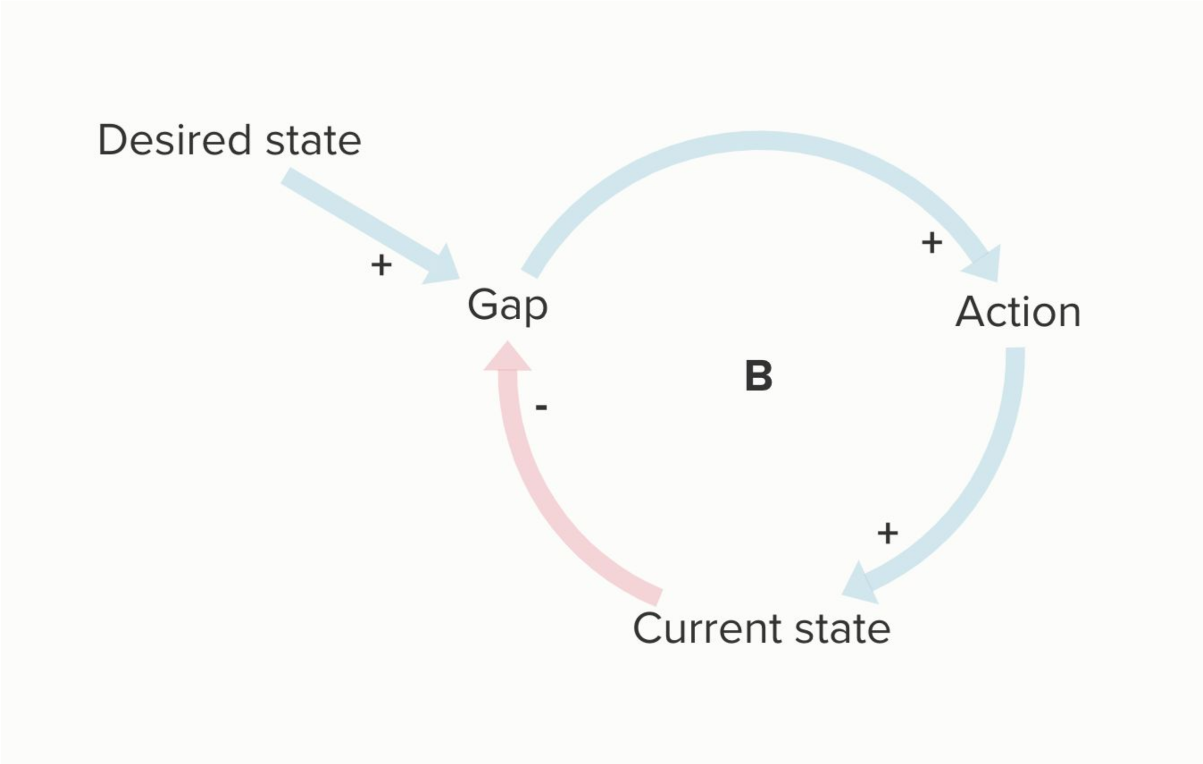

Balancing Feedback Loops (B)

For example, when the Temperature in an apartment is lower than the Desired temperature, then there is a Temperature gap. The heating system will sense this (via a thermostat) and will generate Heat such that the Temperature gap is eventually reduced to zero (i.e, such that the Temperature equals the Desired temperature). If the Temperature drops below the Desired temperature again, then more Heat will be generated until they are again equal. This is a balancing loop.

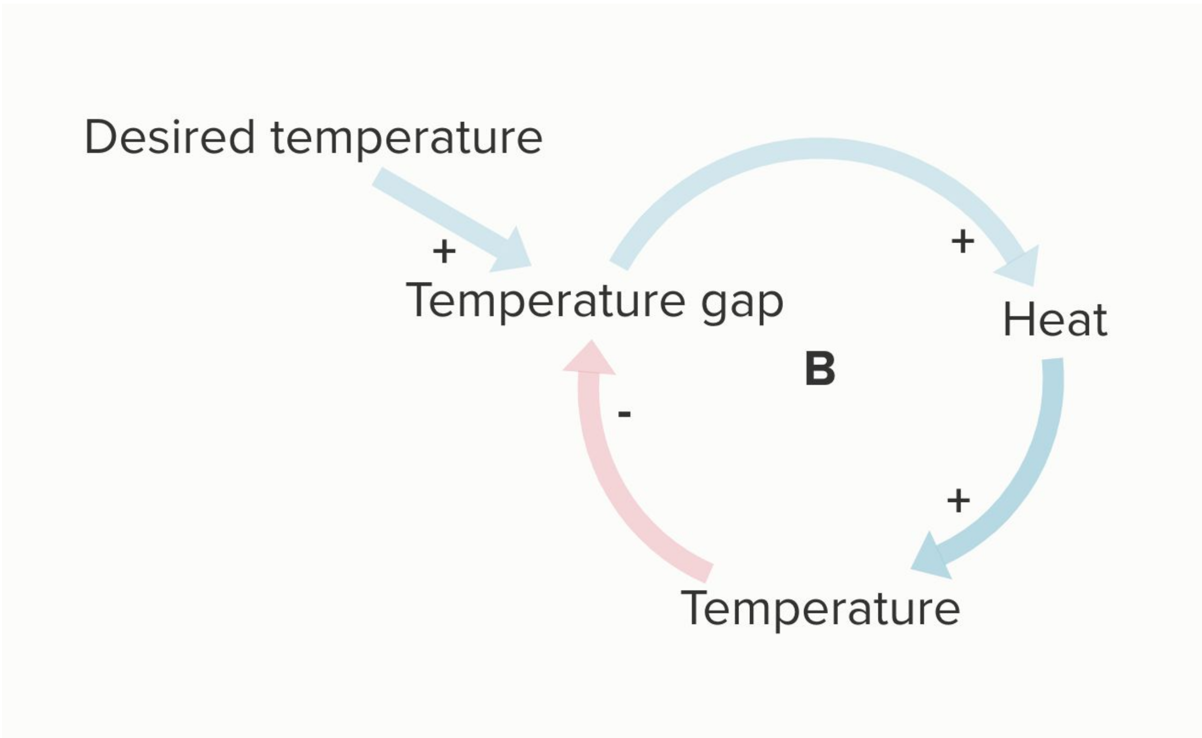

Feedback loops (both reinforcing and balancing) can be found in all walks of life, from psychology to ecology. Before moving on to the next section, we recommend that you watch the excellent video linked below, which should help you get a better grasp of how important and how pervasive feedback loops are in our lives and our planet.

End of Section 3.2: Feedback Loops

That concludes Section 3.2. We hope that by now you:

- Understand that there are two different types of feedback loops: reinforcing and balancing
- Can determine whether a feedback loop is reinforcing or balancing
- Understand that reinforcing loops promote change in a system, whereas balancing loops negate change/promote stability
- Understand that the way in which goals are achieved can often be represented by balancing loops
- Begin to recognise feedback loops around you, in nature, society, the economy, and even in your life in general

In the next section we will look at another tool for helping us understand systems: behaviour over time graphs.

Section 3.3: Behaviour Over Time Graphs

Now that we have an understanding of what reinforcing and balancing feedback loops are, its time to start thinking more deeply about what effects they can have on a system. One of the best ways to do this is to draw or visualise behaviour over time graphs. These are graphs over time of the most important variable(s) in a model. They demonstrate how the variable(s) is likely to develop over time, given the system structure that it is part of. If we refer back to the Iceberg analogy, a behaviour over time graph essentially illustrates the system on the level of patterns.

Exercise 3F

Lets look again at the system structure between Interest earned and Money in bank account.

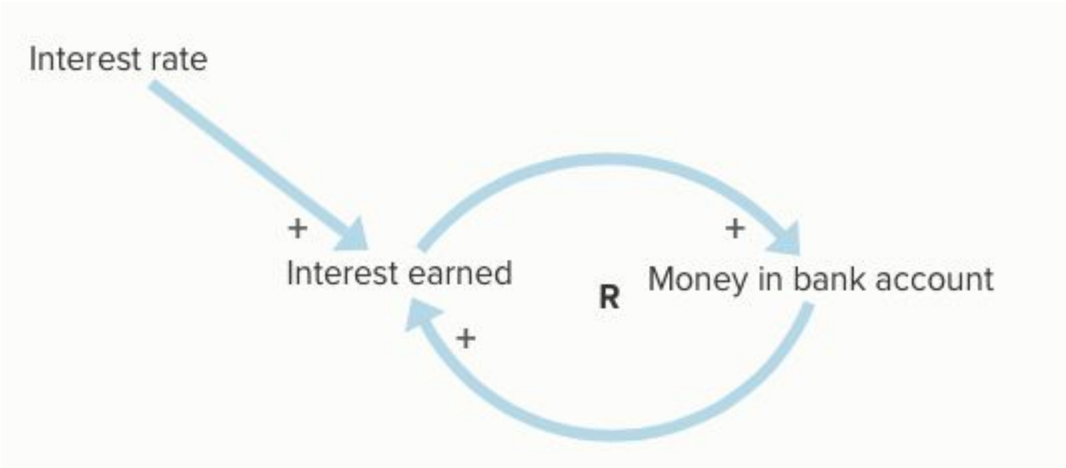

What do you think happens to the Money in bank account over time? Assume that it starts as 1000 USD in year 0, and that the Interest rate is 0.4 and is compounded annually (hint: this means that the Money in bank account is doubled approximately every 2 years - this is an unrealistically high interest rate but is used for illustrative purposes!).

How large will the Money in bank account be twenty years from now, assuming it grows unhindered? Make a quick guess (no need to get out a calculator) and sketch it on a piece of paper with a graph similar to what is seen below:

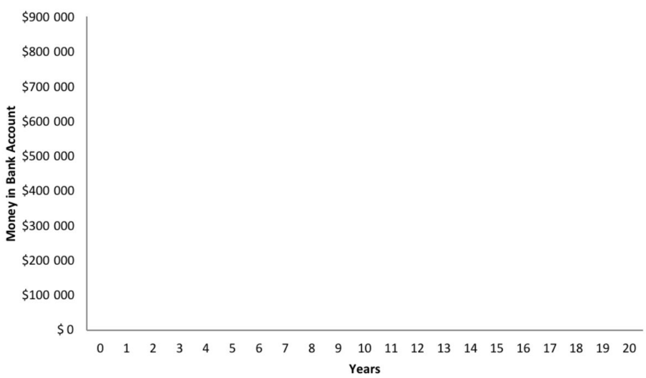

Answer to 3F

The graph below shows the behaviour of the Money in bank account over time (don't worry, we most definitely used a calculator to get this!). In twenty years the Money in bank account grew from $1000 to almost $900,000, without any deposits from the owner!

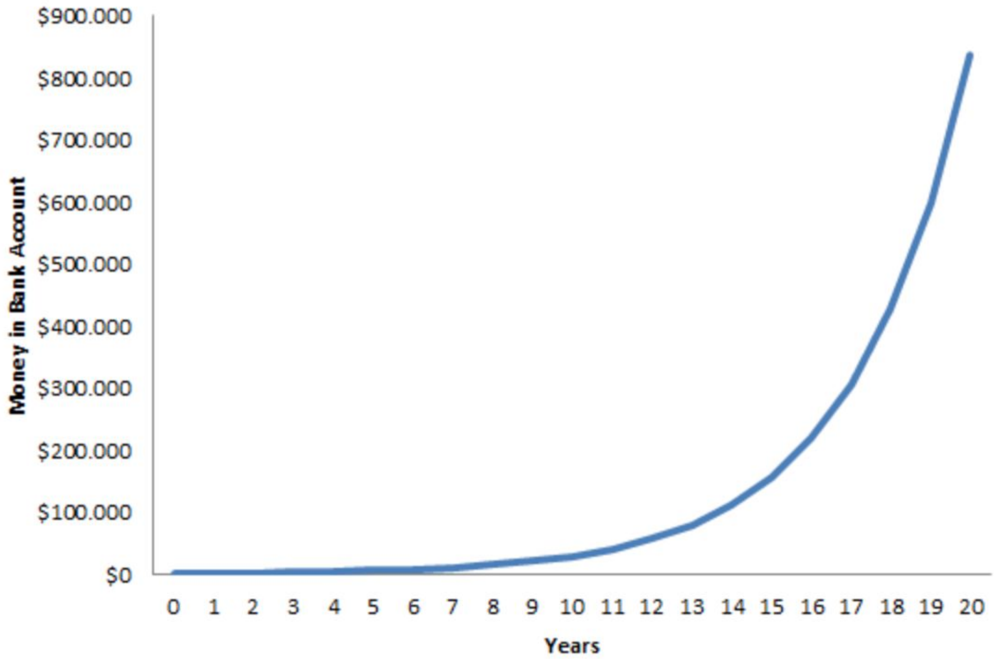

This behaviour is called exponential growth. Exponential growth is created by reinforcing feedback loops as the change of a variable is amplified over time. The more Money in bank account, the more Interest earned, and the more Interest earned, the more Money in bank account!

Many people underestimate the power of exponential growth. To see a more visual illustration of its power, imagine that you were to put one grain of rice on the first square of a chess board, two on the second, four on the third, eight on the fourth, following this pattern all the way up to the 64th square on the board. How much rice do you think you would need to do this?

Other Examples of Exponential Growth

The growth of a bank balance through interest payments is just one example of exponential growth driven by a reinforcing feedback loop. There are many others. For example, as we have seen, a population grows via the reinforcing feedback loop between births and the population (the more births, the more people, the more people, the more births). This goes a long way to explaining why the world population's historical growth looks the way it does, exhibiting a clear pattern of exponential growth.

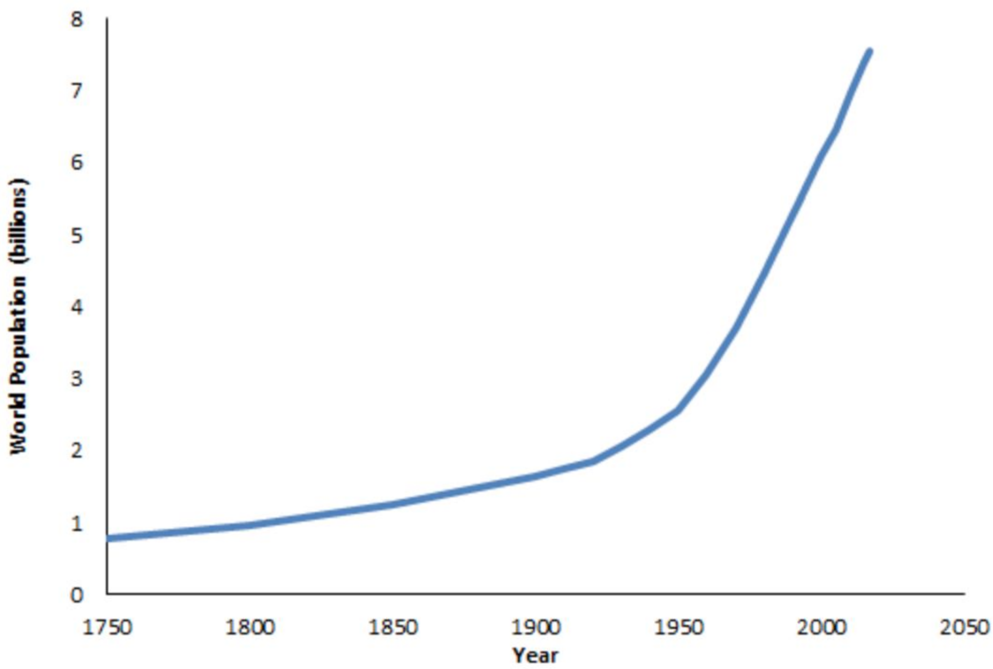

This also shows us how a system's structure is what produces its behaviour. Although the money in a bank account and the global population seem like unrelated systems, they have similar structures (i.e. they are dominated by reinforcing feedback loops) which means that they exhibit similar patterns of behaviour, which in this case is exponential growth.

Goal-seeking Behaviour

The structure of balancing feedback loops produce a pattern of behaviour which we call "goal-seeking behaviour". To demonstrate what this means, let's look again at the thermostat example:

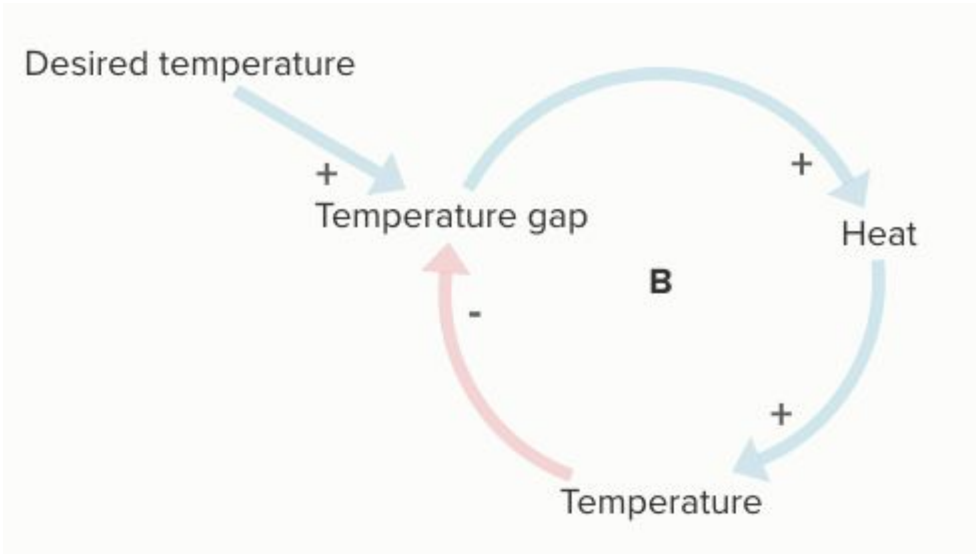

Suppose the Temperature in the apartment is 5 °C degrees while the Desired temperature is 25 °C. The Temperate gap between the desired and actual temperature will cause thermostat to generate Heat such that the Temperature will gradually approach its goal of 25 °C. Initially, the gap between the Temperature and the Desired temperature is large and as such the heating system will work full power to make up for this gap. However, as the Temperature gap gets smaller, the heating system will not need to provide as much Heat. As seen in the behaviour over time graph below, the increase in Temperature is significant at the beginning (as the heating system works full power) and gradually becomes less and less as the Temperature reaches the goal of 25°C.

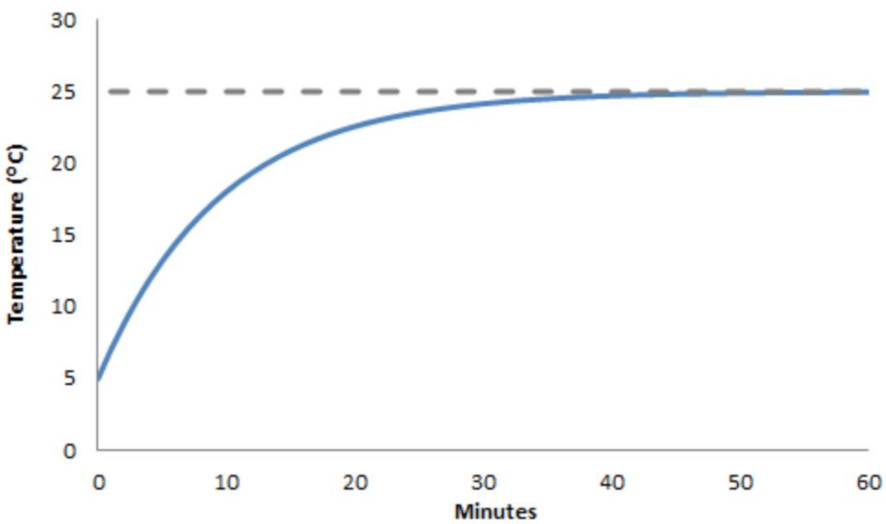

Oscillation

When there is a delayed cause and effect relationship in a balancing loop, this can result in what is known as oscillating behaviour.

You may have experienced oscillation whilst in the shower. When the water is too cold, you might open the hot tap more. However this doesn't cause the water to heat up immediately, as the heating system takes time to react to your adjustment. This might cause you to impatiently open the hot tap even more, such that when the water does eventually arrive, it is far too hot! This might then cause you to quickly turn the hot tap way down, with the result that the water eventually arrives as too cold!

This structure of this system can be represented as seen below. Note that the double dashed arrow from Shower faucet setting to Water temperature represents the fact that there is a significant time delay between the adjustment of the faucet setting, and the change in the water temperature. In other words, the double dash on an arrow means that there is a significant delay between the cause and effect.

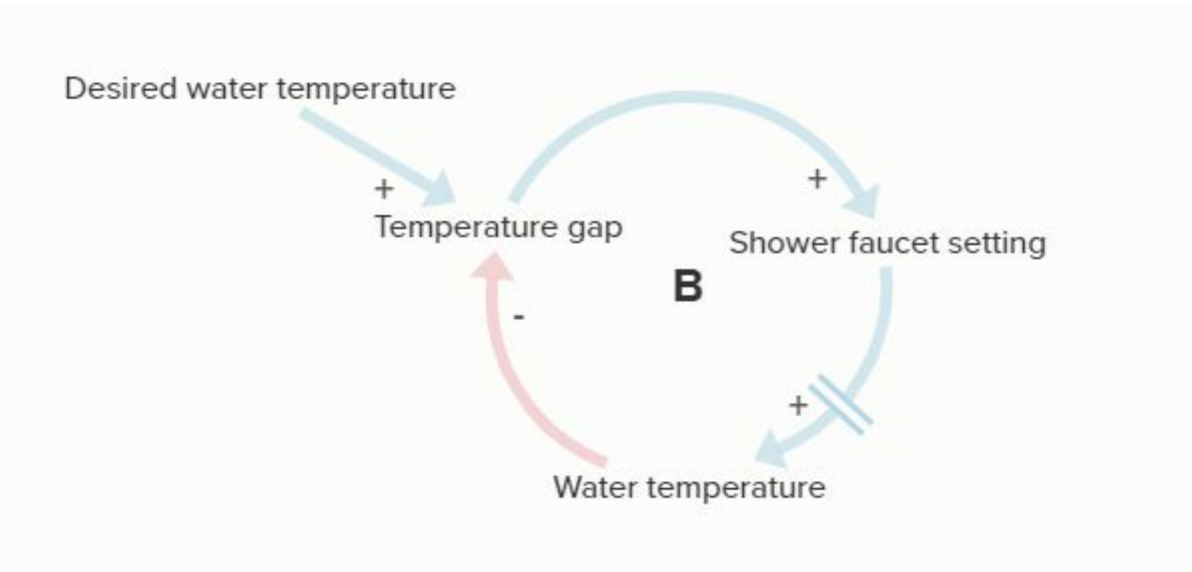

Below we can see a behaviour over time graph that depicts how the temperature of the water might oscillate around the desired temperature (the dotted line), based on this system structure.

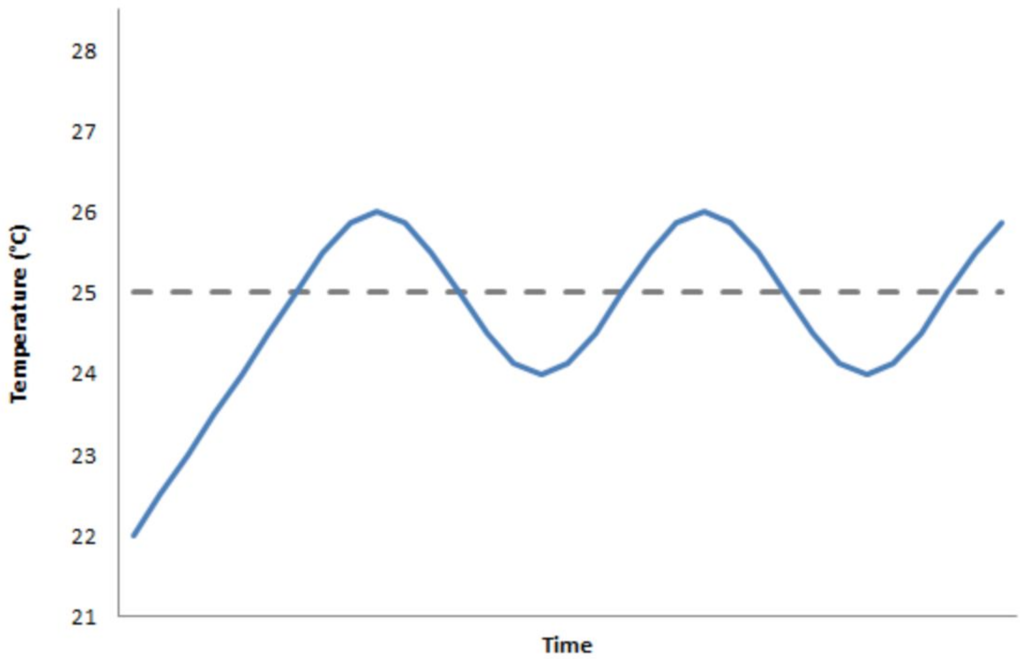

S-shaped Growth

Many if not most systems are made up of both balancing and reinforcing feedback loops. The existence of these two kinds of loops in a system can result in what is known as "S-shaped growth," which can be seen in the graph below.

A variable exhibits S-shaped growth when it is at first dominated by a reinforcing feedback loop, and then later dominated by a balancing loop as it reaches some kind of limit to growth. Systems thinkers refer to this situation as a shift in loop dominance in the system.

For example, consider how a population of deer might initially grow exponentially (dominated by the feedback loop between births and population), and then, as the population grows towards the carrying capacity of its area, the death rate of the deer might increase (or the birth rate might decrease), such that the death rate equals the birth rate, and the population of deer reach what is known as a dynamic equilibrium (i.e., a steady state).

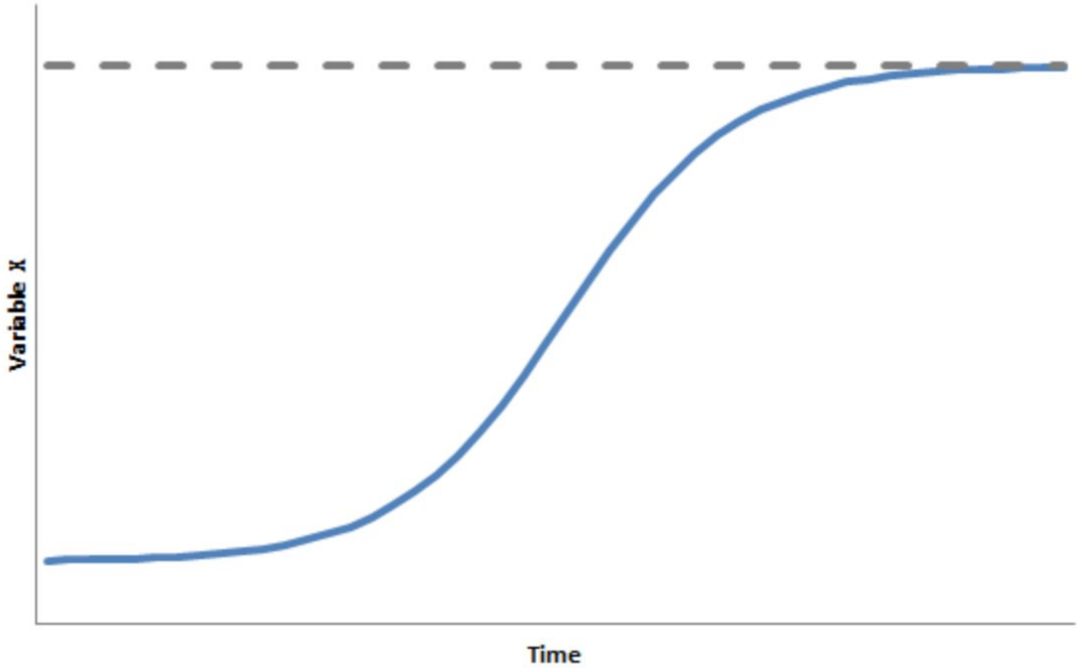

Exercise 3G: Feedback Loops

Take a look at the CLD on the below. It is a generic model of how a firm's Marketing efforts affect their Sales.

- Q1: Can you put the correct signs next to the arrows?
- Q2: Can you identify and define the type of feedback loops present?
- Bonus Question: Once you have answered Q1 and Q2, can you sketch the behaviour of the firms Sales on a graph over time? Assume that the Market saturation is at first low (imagine the firm has just launched a new product), but eventually grows to full saturation over time.

Answers to 3G (1/3)

In this system there is a reinforcing and a balancing loop. The reinforcing loop says that an increase in Marketing efforts increases Sales, and as Sales and Revenue continue to grow, the company has the ability to further increase its Marketing efforts.

If we had a mental model that only recognized this loop, then we would presume that increasing Marketing efforts always led to increases in Sales.

Tip: to determine if a loop is reinforcing or balancing you can:

1. Follow the loop and the logic of each relationship. For example, start at one variable and imagine what will happen if you increase that variable. If you follow the loop around and an original increase in that variable leads to another increase, then the loop is reinforcing. If it leads to a decrease in the variable, then the loop is balancing;

OR

2. Count the positive and negative signs. A balancing loop always has an uneven number of negative signs (at least one), whilst a reinforcing loop has either no negative signs, or an even number of them.

Answers to 3G (2/3)

However when we look at the wider system in which Marketing efforts operate, we see that increased marketing will only lead to increased Sales for so long, because when Sales increase, then the level of Market saturation also increases. This is because everyone who would potentially buy the product (i.e. the addressable market) is already doing so. This highights a balancing loop in the system.

At some point the market will reach total saturation. Increasing the Marketing efforts then has less effect on Sales (it does not cause Sales to increase, as before, but rather just keeps the same amount of people buying the product). This may seem strange to the management team, as increasing their Marketing efforts previously always led to more Sales.

Of course many companies who have huge sales still engage in marketing. And this model is not saying that they are foolish to do so. The likes of Apple and Microsoft still need to advertise their products if they want to keep the customers they have won. The point this model is trying to make is that marketing executives at growing companies should not be totally surprised when the effectiveness of their marketing efforts (in terms of generating an increase in sales) begins to decline as their product reaches market saturation.

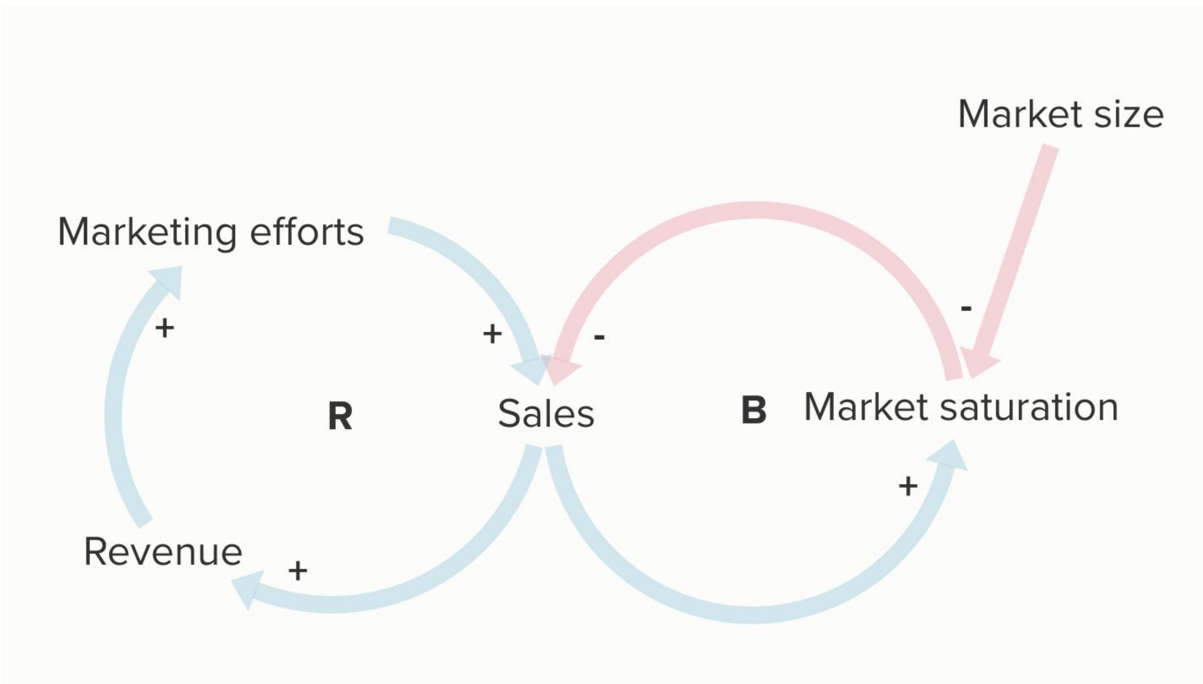

Answers to 3G (3/3: Bonus)

id you estimate that Sales would develop in the way shown on the graph below? Our thought process behind reaching this answer went something like this:

"Initially, there is no market saturation, which means the balancing loop is not so strong. This means the reinforcing loop is likely to dominate, resulting in an exponential growth in sales. However, growth in Sales eventually leads to significant Market saturation, such that Sales slow down and finally stabilizes when full Market saturation (shown by the dotted line on the graph) is reached (i.e., when the entire addressable market already owns or is consistently buying the product)."

As such, the system shows S-shaped behaviour, which is typical for systems with a combination of a reinforcing and a balancing feedback loop.

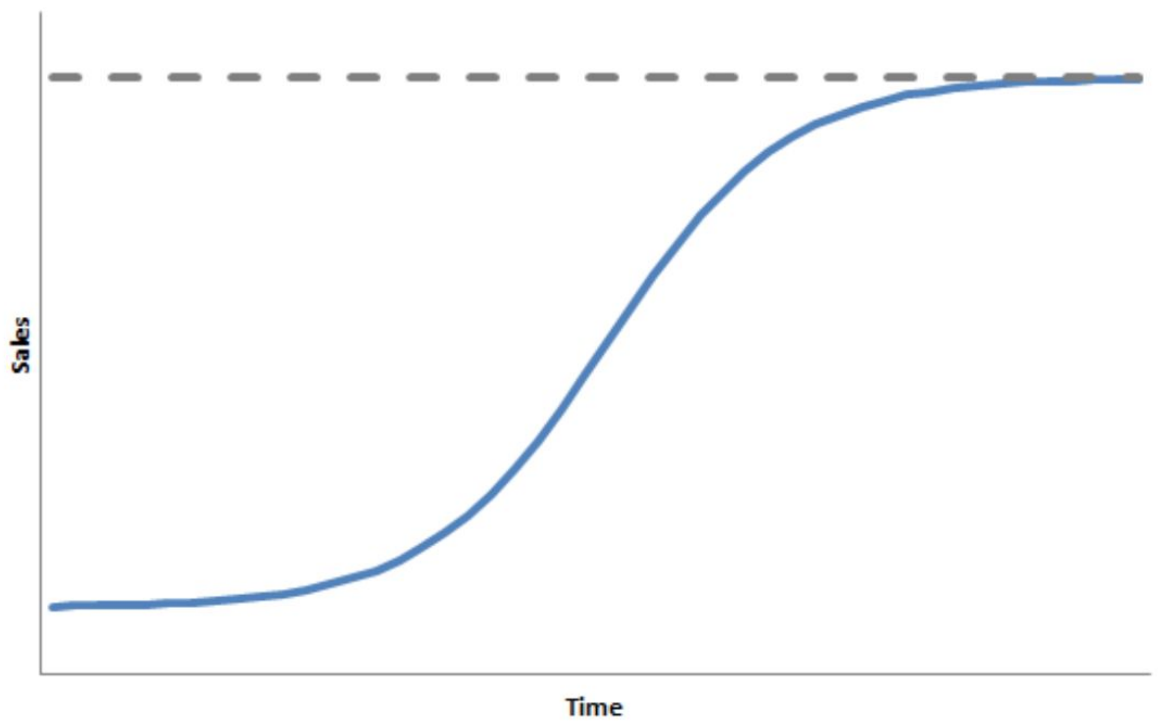

End of Section 3.3: Behaviour Over Time Graphs

That concludes section 3. We hope that by now you:

- Realise the value of behaviour over time graphs in understanding how a systems structure leads to its behaviour
- Understand the power of exponential growth, which is typically driven by a reinforcing feedback loop
- Understand that balancing loops often produce goal seeking behaviour
- Understand that delays in balancing loops can often result in oscillation around a goal
- Understand that a combination of a reinforcing and a balancing loop often leads to S-shaped growth

Section 3.4: Systems Archetypes

In this section we will look at what are known as systems archetypes.

Limits to Growth: A Systems Archetype

In the scenario of the previous exercise, more Efforts initially led to improvements in Performance (measured by sales) (R). However, over time that Performance reached a limit or Constraint that caused the Performance to slow down (B).

Recognising this scenario leads to the understanding that you should either look for ways to remove the limits or else prepare to adapt to them. For example, the company in the previous exercise could have explored new markets to increase its addressable market (removing the constraints of market saturation), or could have planned (e.g. budgeted) for the fact that marketing efforts were unlikely to result in the same growth in sales as was previously experienced.

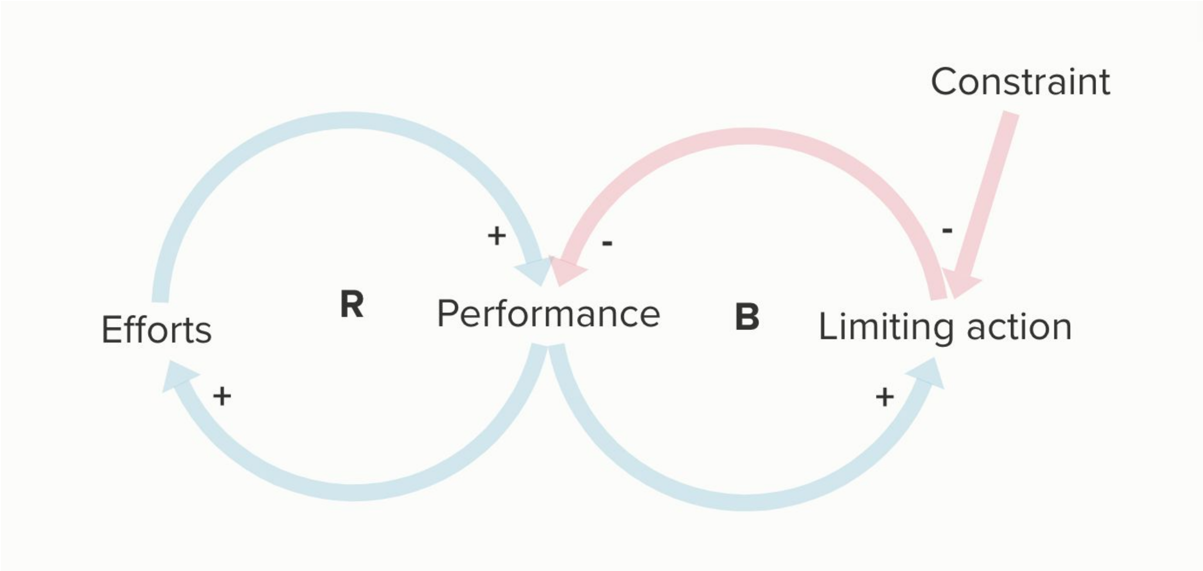

Archetypes

The scenario in the previous page is an example of a systems archetype (Limits to Growth). Systems archetypes are combinations of reinforcing and balancing feedback loops that commonly occur across different systems.

Each archetype has a characteristic theme and storyline, behavioural pattern, and potential for action. Being able to recognize system archetypes in different situations allows for a deeper and faster understanding of that system, and could also help us design effective interventions. By applying archetypes habitually, you will learn to “see” structures when you hear stories that are similar to the archetype.
Let's look at some more examples.

Tragedy of the Commons Archetype

Imagine a small lake in rural Vietnam where everyone is allowed to fish. The lake doesn't belong to anyone in particular and you can fish as much as you want. Imagine that there are two fishing companies in the area and that all those who fish in the lake belong to either company A or company B. As can be seen on the right, the more fish company A catches, the more profits they make and the more they will increase their Fishing (R1), as profits allow the company to hire more fishermen and buy more fishing boats. This principle is the same for Company B (R2) and together both companies make up the Total fishing.

In the beginning both Company A and Company B make good profits and provide a valuable service to the community.

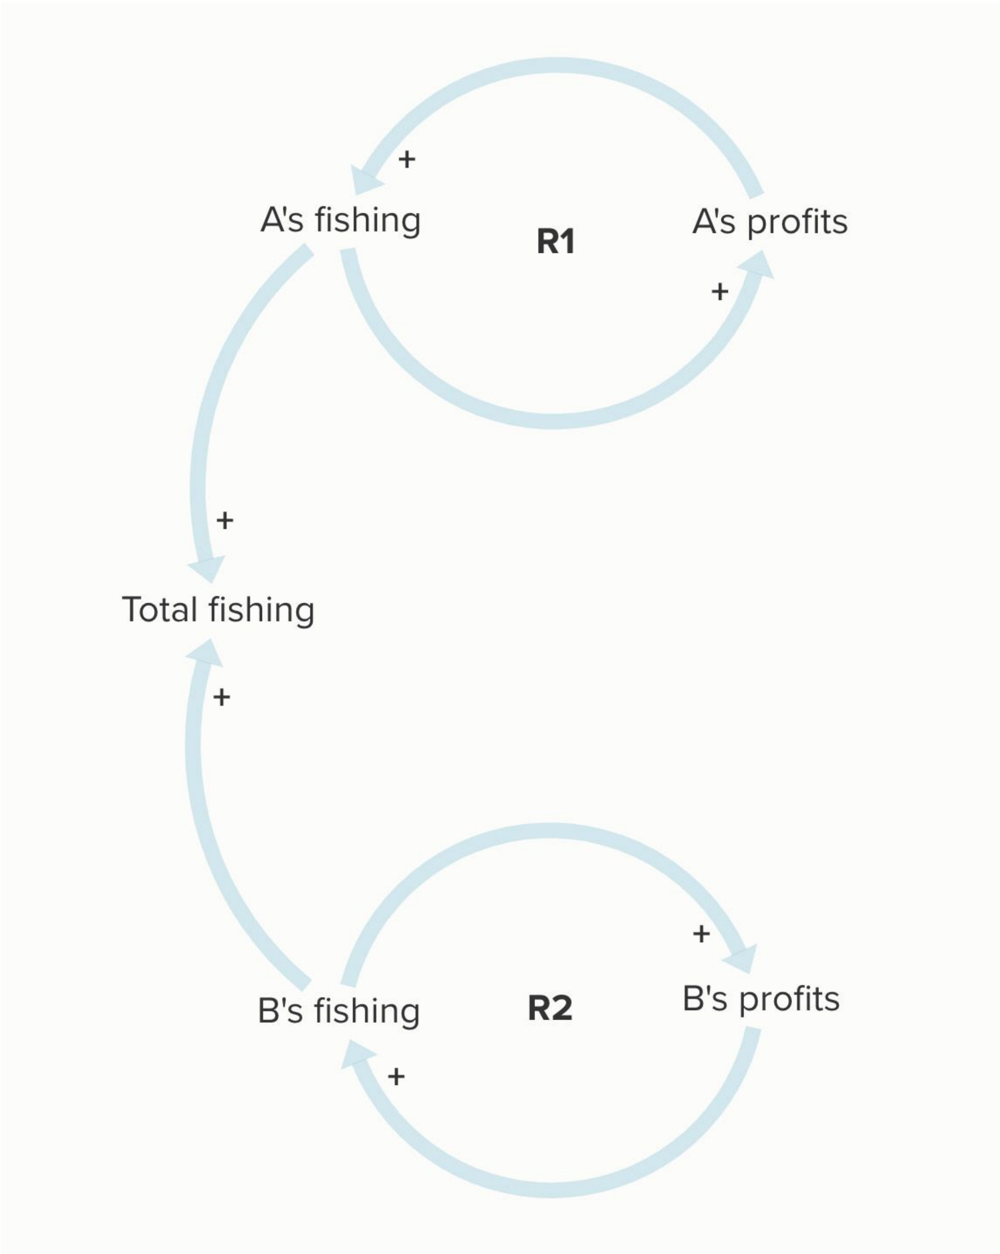

However, when we look at the longer term we see that this system, if left uncontrolled, may result in some unintended consequences. If the level of fishing is not limited in some way, then it may eventually lead to a situation whereby more fish are being caught than are being naturally replenished, leading to a decline in the number of fish in the lake.

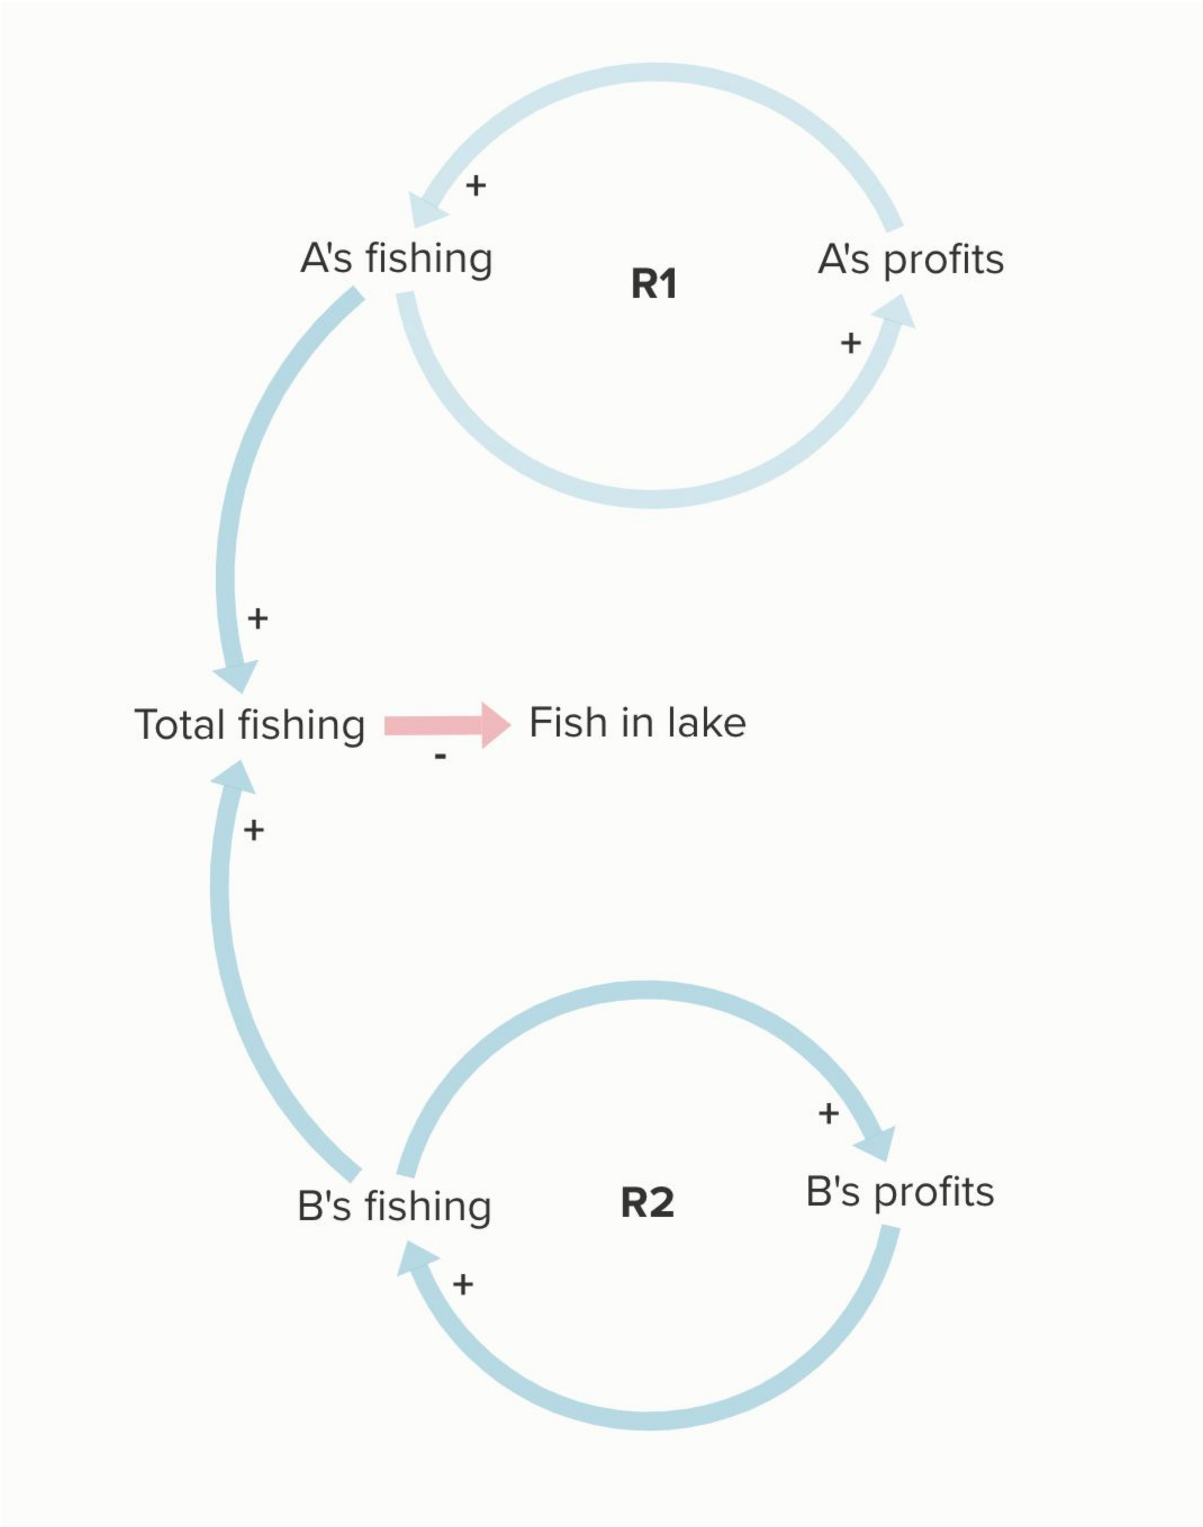

Less fish in the lake in turn leads to more effort to catch a fish, as fishermen might need either better boats to go to deeper waters, or to spend more time fishing to catch the same amount of fish. This reduces the profits of both companies, as well as the fish supplied to the community. It also damages the fish population and perhaps the whole aquatic ecosystem.

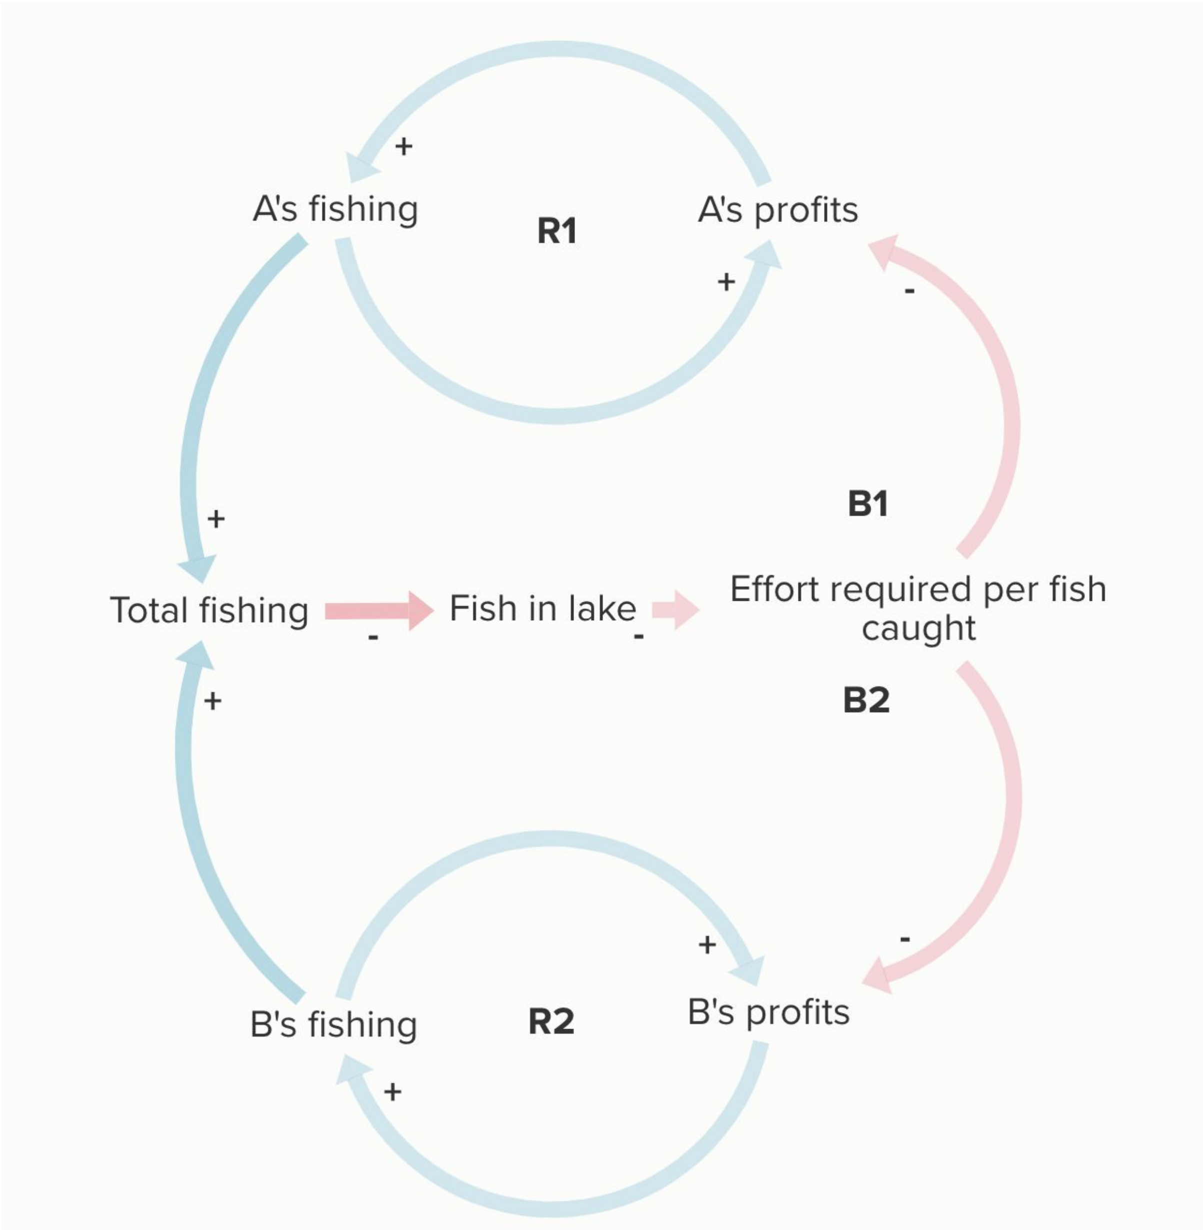

The scenario in the previous example describes the "Tragedy of the Commons" archetype and is often referred to when people discuss environmental issues (the phenomenon became widely known through an article by ecologist Garrett Hardin in 1968). In the Tragedy of the Commons archetype, individual actors each behave rationally according to their own, short-term self-interest (R1 and R2) but that eventually results in the depletion or spoiling of a shared resource (B1 and B2), which is bad news for all involved.

Possible solutions to Tragedy of the Commons type problems are to:

1. Educate actors about the consequences for the population as a whole
2. Reduce the activities by regulating access to the resource or privatizing the resource so that the owner take longer term considerations into account.
3. Reduce the information delay between the total activity and the resource being depleted.

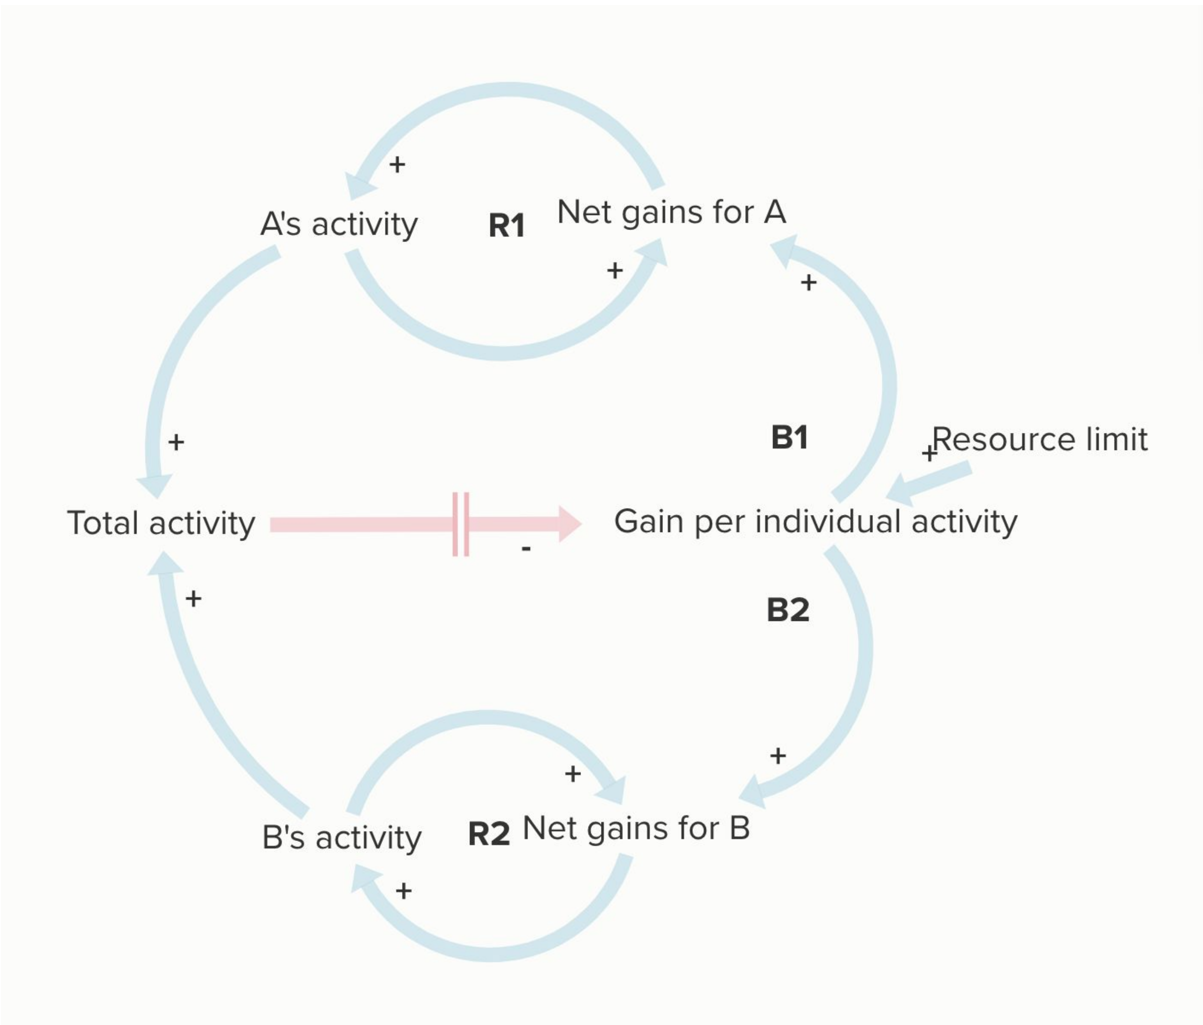

The below video gives a more in depth and colourful description of the Tragedy of the Commons. If you have the time, we suggest that you watch it!

Success to the Successful Archetype

On the right you can see another archetype, known as the "Success to the Successful" archetype. This archetype explains how one actor in a system can become far more successful than other actors due to small differences in initial conditions, expectations, or random factors.

For example, when the coach of a football team thinks that player A is better than player B, he or she will spend more time training player A. The time of the coach is limited, so player B will receive less training than A. Because player A receives more training/attention, he or she might improve even more than player B, further widening the gap in the skills between them. This will further justify the coach's decision to train player A instead of player B.

This self-fulfilling prophecy continues to reinforce the success of player A while diminishing the success of player B, all because of an initially small difference in skill between them.

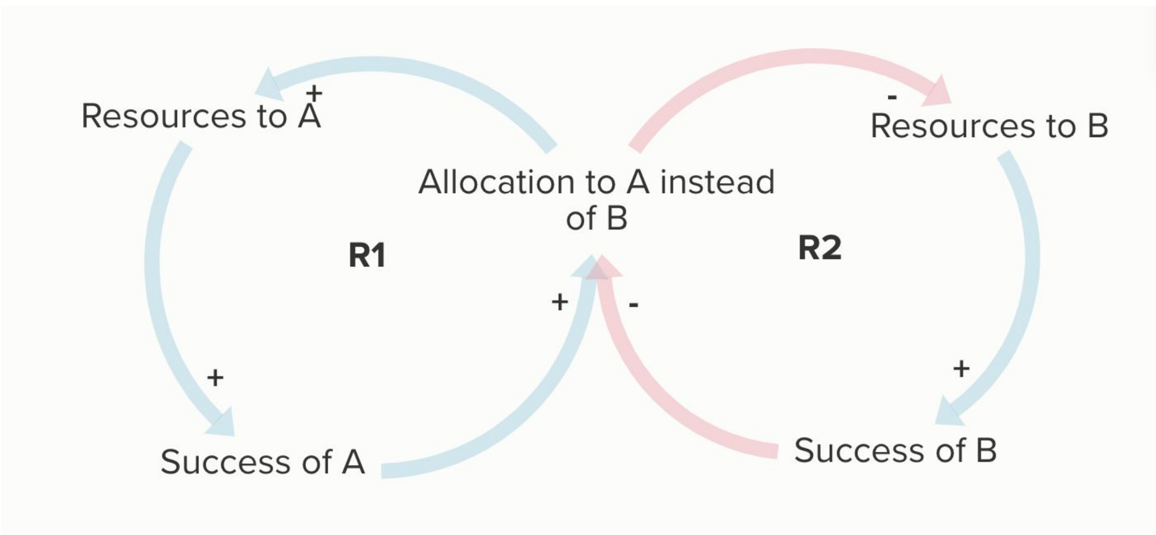

This archetype applies to a variety of situations and can explain for example:

- Why there is social and economic inequality within and between many countries
- Why the rich are getting richer and the poor are getting poorer
- How one species can drive another to extinction
- Why success at school leads to success later in life

There are two main strategies for dealing with similar situations:

1. Redistribute the unequally distributed resource. An example of this is progressive tax scales or property taxes.
2. Separate the two reinforcing loops so that the two entities are not competing for the same resource anymore. For example, if a fishing community who was previously in competition with another began fishing in a new lake.

Exercise 3H: Archetypes

We encourage you to watch the three videos below. Can you identify which systems archetype is applicable to each of the news stories? The answers are on the next page.

1. Delhi chokes on air pollution

2. End of the outbreak: Animated look at the ebola epidemic

3. What discrimination looks like in America

Answers to 3H

The "Tragedy of the Commons" archetype is applicable to the first video. In this case, the shared common resource is clean air. This essential resource is being spoiled by individual actors that behave rationally according to their own self-interest (by driving cars or burning crops, for example), despite the fact that spoiling the resource is contrary to each others' long-term best interest (public health).

The Limits to Growth archetype is applicable to the second video. In this case, the virus (Ebola) initially spread rapidly but was eventually slowed down when there were less and less people to infect, as people were either already infected or those infected were put into quarantine/isolation.

The Success to the Successful archetype is applicable to the third video. Due to initial prejudices and discrimination, a minority group (black people in the US) became socially disadvantaged. This social disadvantage was then interpreted not as the result of earlier prejudice and discrimination but as evidence that the minority is indeed inferior, which resulted again in prejudice and discrimination.

End of Section 3.4: Systems Archetypes

As you can see, archetypes can be used to identify the systemic structure of a problem and to remind us about potential high leverage solutions to those problems. When you become familiar with the different systems archetypes, you might start seeing them everywhere. You might understand humankind's history through the lens of the Success to the Successful archetype, or you might see a messy kitchen as the result of a Tragedy of the Commons-type structure between yourself and your housemates!

Being aware of archetypes can help us to identify the systemic structure of a problem and to remind us about potential high leverage solutions to those problems. Relating a problem to an archetype can help you to see how a problem might be similar to many others, and can help you to see common solutions to those problems. For example, you might recognize that the collapse of fish stocks and global warming might both be results of a "Tragedy of the Commons"-type structure, and that both might have many common solutions, such as increasing awareness about the problem, introducing legislation, improving technologies, and much more.

To learn more about archetypes and see examples of other archetypes, we recommend the work of Peter Senge (e.g., The Fifth Discipline, 1990), Donella Meadows (Thinking in Systems, 2004: chapter 5) and this article written by Daniel Kim. For more readings on leverage points in systems, we recommend this article by Donella Meadows.

End of Module Three

This is the end of Module Three. By now we should have a basic understanding of CLDs and their uses. To recap, we have studied the following topics in this module:

- Causal relationships (positive and negative) between variables
- Different kinds of feedback loops (reinforcing and balancing)
- Behaviour over time graphs
- Archetypes and leverage points

In the next module we will look at more examples and you will further develop your skills in building your own CLDs.

Further Readings

- Harari, Y.N. (2015). Sapiens : a Brief History of Humankind. New York: Harper.
- Hardin, G. (1968). Tragedy of the Commons. Science 162 (3859), 1243-1248.
- Kim, D.H. (1992). Systems Archetypes I. The Toolbox Reprint Series, Pegasus Communications, Inc.
- Meadows, D.H. (2008). Thinking in Systems. White River Junction, VT: Chelsea Green Publishing
- Senge, P. (1990). The Fifth Discipline.
- New York: Currency Doubleday

Module Four
Building Your Own Systems Models

WASSILY KANDINSKI [CC0 OR PUBLIC DOMAIN], FROM WIKIMEDIA COMMONS

Module Four Overview

Introduction

In the previous module, we learned how to construct and read CLDs, as well as how to understand some of the phenomena they can help us to recognise, such as feedback loops and systems archetypes.

However, the process of building your own models from scratch is another challenge, and that is what we will focus on in this module. To do so, we will model different problems and describe our thinking throughout the process.

There is always more than one way to model a system, and our process for doing so is not necessarily the best or only way. Nonetheless, we hope that by walking you through our thought process for building models, you might learn some good tips for your own model building. We also hope that the models in this chapter will provoke some thoughts about some interesting topics!

Module Goals

By the end of this module you should be able to understand:

- How to develop CLD models from scratch
- How to analyse and infer useful insights from CLDs

We will be asking you to develop CLDs for three different topics, and we will provide you with a certain amount of information for each topic. We will start with a model about a fairly simple and relatable topic, and then move on to more complex topics.

Module Outline

This module is divided into six sections, each of which is described below:

1. Modelling Procrastination and Productivity

2. Modelling the Utility Death Spiral

3. Modelling the Theory of the Invisible Hand

4. A First Experience of Group Model Building

5. Tips for Model Building

6. Identifying Stocks and Flows

Section 4.1: Modelling Procrastination and Productivity

As mentioned in Module Two, CLDs can be used to represent different perspectives, or mental models, about a certain issue or topic. In this section we will create two different CLDs to represent two different mental models about the same topic. The topic we will focus on is how procrastination affects productivity.

Let's imagine that you interview two different university students about the topic, Mary and Paul. In response to your question about how procrastination affects productivity, Mary says the following:

Mary's Mental Model: Procrastination can be useful for productivity

"When I am facing a deadline for an assignment, I usually wait until the last minute to start working. I used to think that this was a terrible habit, but over the years I've come to realise that I am simply more productive when under pressure. This is because the closer it gets to my deadline, and the more remaining tasks there are to complete, the more deadline pressure I feel. That deadline pressure increases my productivity, such that I complete more tasks per hour. This means that overall I spend less time on the assignments, with the same outcome, which I would consider to be a more productive situation!"

Imagine that you then turn to Paul, and he answers your question as follows:

Paul's Mental Model: Procrastination can increase errors and lower productivity

"I share Mary's perspective that when I'm under pressure I complete tasks quicker, but I would add that when I'm working faster because of time pressure, I make more errors. And more errors means more remaining tasks to complete. Also, if I don't spot the errors, then this will reduce the quality of my work. As such I think I am usually better off when I start the assignment early and complete it with less pressure and thus less errors, rather than waiting for the deadline pressure to kick in."

Attempting a Model From Scratch

In the next slides we will guide you through a process for developing CLDs of Mary and Pauls' mental models. However, before we begin, we encourage you to make your own attempt at developing a CLD of each mental model. You can then compare your models to ours and see if you have learnt anything from our process. Please note that model development is a challenging process that gets easier with practice. You are certainly not expected to be able to produce models quickly and effectively just yet. Nonetheless we believe that the best way to learn is to try!

Exercise 4A: Identify the Variables

We will begin by creating a CLD of Mary's mental model, shared again below:

"When I am facing a deadline for an assignment, I usually wait until the last minute to start working. I used to think that this was a terrible habit, but over the years I've come to realise that I am simply more productive when under pressure. This is because the closer it gets to my deadline, and the more remaining tasks there are to complete, the more deadline pressure I feel. That deadline pressure increases my productivity, such that I complete more tasks per hour. This means that overall I spend less time on the assignments, with the same outcome, which I would consider to be a more productive situation!"

As a first step in modelling, it can be useful to pick out some of the essential variables that you will use to build the model. Try to do this yourself before going further. Note that this is often an iterative process and you may choose quite different variable names when you begin drawing causal connections. Oftentimes, many modellers will simply come up with the necessary variables as they build the model. However we think deciding on the variable names before beginning the model can be a useful exercise for a beginner.

Here are some helpful tips for choosing variables:

- Choose only those variables that are necessary to communicate the essential message of the model
- Look for relationships and causality, and let them guide you to pick up important variables. Relationships will involve two or more variables
- Do not include verbs in your variable name; verbs are actions and they are represented by the causal arrows
- Use clear nouns to represent variable names
- Variables are usually referred to in the text multiple times with different names; try to group them into one variable name.

Answer to 4A

Below you can see the four variable names that we chose, along with some viable alternatives for that variable, and some names which we would not consider best practice, for the reasons described.

1 *Remaining tasks*

A very important variable that you can identify is the number of *Remaining tasks*, which we consider to be the total number of tasks that Mary must compete before she considers her assignment finished. This includes tasks such as reading papers and writing chapters/sections of her paper.

Viable alternative names: *Open tasks* / *Mary's remaining tasks* / *Uncompleted tasks* / *Tasks* / *Fraction of assignment left to complete*

Not best practice: *Lower amount of tasks* would not be an advisable variable name because causal connections will already allow us to think about how these variables change over time. Remember that variable names should be nouns whereas verbs are taken care of by the arrows!

2 *Days remaining to finish tasks*

This variable can be used to measure the amount of days Mary has until the tasks deadline.

Viable alternative names: *Time left* / *Remaining time*. You could have three variables to represent this one, where you have *Task deadline date* and *Current date* both determining the *Difference between tasks deadline and current date*, which is the same as the *Days remaining to finish tasks*. However we believe that *Days remaining to finish tasks* captures the meaning in a simpler way.

Not best practice: *Closeness to deadline* could be used as a variable, but we believe it is not as clear a term as *Days remaining to finish tasks*.

3 ***Deadline pressure***

Mary has spoken of how the pressure to meet her deadlines plays a crucial role in motivating her to finish tasks quicker. As such it's very important to represent this variable.

Viable alternative names: *Schedule pressure / Pressure to finish assignment*

Not best practice: *Lack of deadline pressure*. This would not be an advisable variable name because it is easier to think of increasing deadline pressure than it is to think about decreasing lack of deadline pressure!

4 ***Task completion rate***

This represents the number of tasks Mary completes per hour/day/week

Viable alternative names: We could have chosen *Productivity* as a variable name. However we believe that *Task completion rate* is a more explicit term that represents exactly what we mean when we talk about productivity.

Not best practice: *Tasks completed*. If we are trying to talk about productivity, it could be confusing if we used the term *Tasks completed* to represent this, because this could mean the total number of tasks *already* completed, which is a result of productivity, and is not productivity itself. Tasks completed *per day*, on the other hand, is a proxy for productivity.

You may have also chosen a variable name such as Procrastination. We did not choose to represent this variable because we think that it is not necessary for our model.

Exercise 4B: Identify the Central Variable(s)

In the previous exercise we have identified four variables that will be used to construct our CLD. Our next step would be to identify the variable that we consider most important, and then move from there to draw the necessary causal relations. Please look at the list below and choose the variable that you believe is most central to our problem.

List of variables:

- Remaining tasks
- Days remaining to finish tasks
- Deadline pressure
- Task completion rate

Answer to 4B

We argue that the Remaining tasks is the central variable to our problem. This is because Mary's main goal is to reduce the number of Remaining tasks to zero — in other words, to finish the assignment.

You may have chosen another variable, such as Task completion rate, as the central variable, arguing that the main issue is productivity. Note that there is no right answer to this question. Indeed many people might just read the text and start developing a model right away without thinking of what is the central variable. We are merely suggesting one approach that we find useful.

Exercise 4C: Identify the Causal Relationships

Once we've identified our central variable, we can move form there to show the cause and effect relationships that it is a part of.
First we can consider the variable(s) affecting our central variable. What variable do you think has an effect on the Remaining tasks?
List of remaining variables:

- Days remaining to finish tasks
- Deadline pressure
- Task completion rate

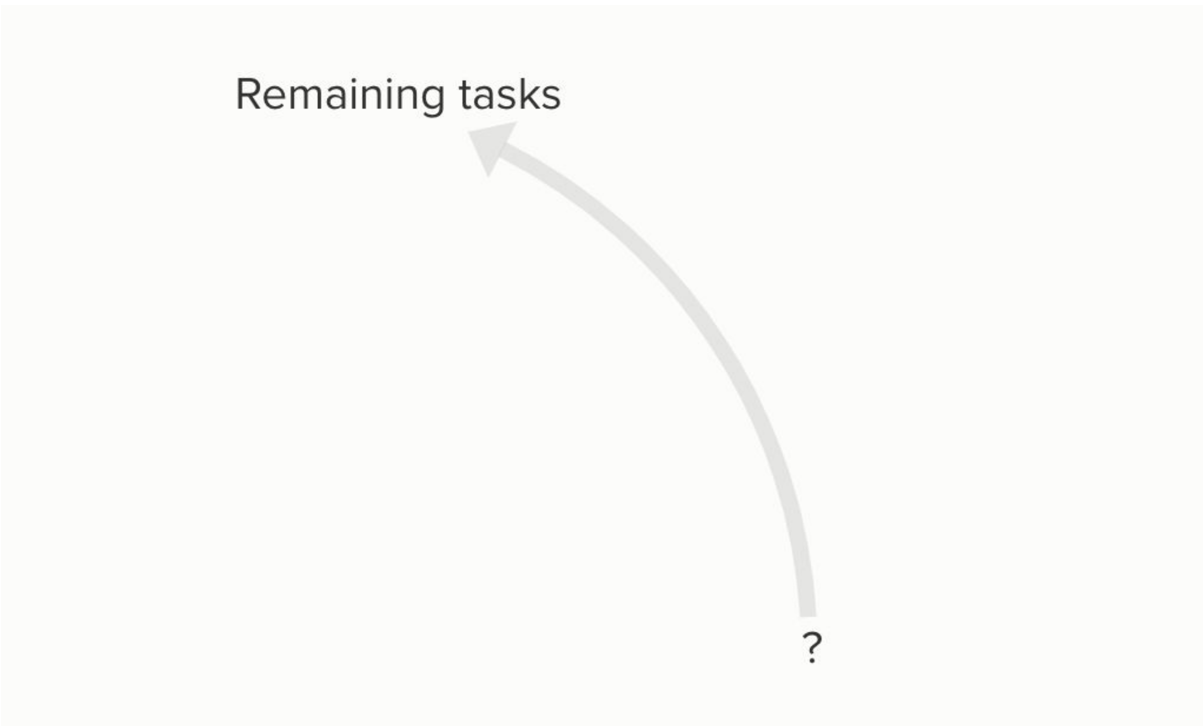

Answer to 4C (1/3)

The more tasks completed per hour/day/week, the less tasks there will be left to do. As such, the Task completion rate affects the Remaining tasks.

What polarity would you give to this relationship - positive or negative? Go to the next page to see the answer.

Answer to 4C (2/3)

An increase in the Task completion rate would result in less Remaining tasks, all else equal. As such, there is a negative relationship from Task completion rate to Remaining tasks.

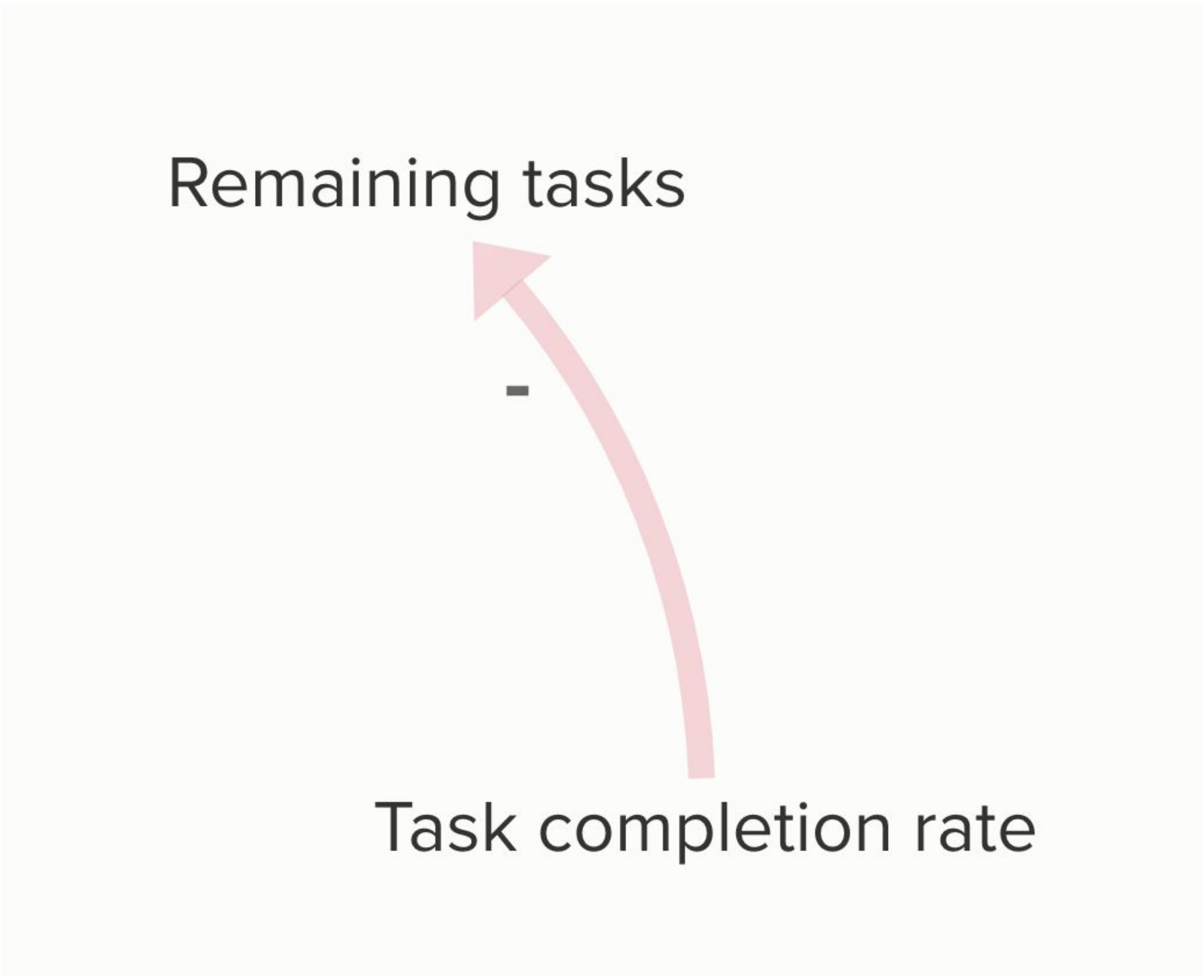

Next we can identify what variable(s) the Remaining tasks is affecting. What do you think this variable is affecting in the model? As a reminder, here is some of what Mary said:

"...the more remaining tasks there are to complete, the more deadline pressure I feel."

See how we model this on the next page.

Answer to 4C (3/3)

Based on what Mary has said, more Remaining tasks means more Deadline pressure. As such there is a positive relationship going from the former to the latter.

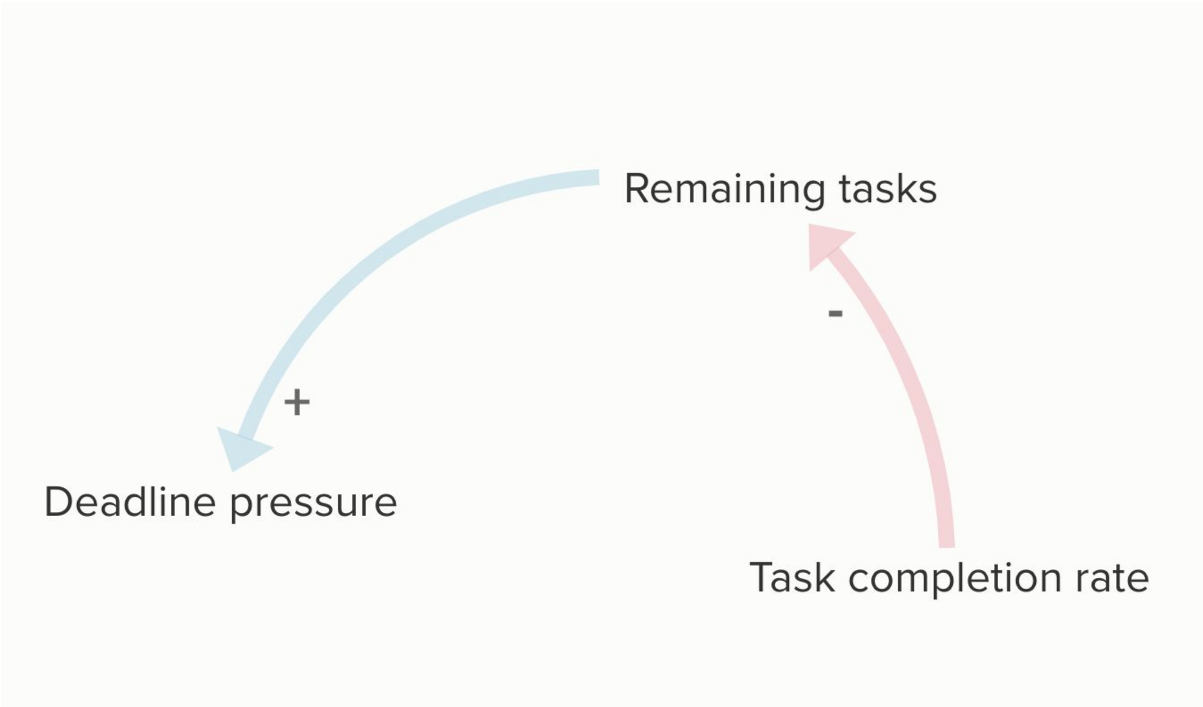

Exercise 4D: Finding More Causal Connections

Our next step would be to identify any other variable(s) that are affecting Deadline pressure.

Mary has said that "the closer it gets to my deadline, and the more remaining tasks there are to complete, the more deadline pressure I feel".

We have already shown how Remaining tasks affects Deadline pressure. However, Mary has also said that the closer it is to her deadline, the more pressure she feels.

How would you model this? Our answer is on the next page.

Answer to 4D

Fewer Days remaining to finish tasks means more Deadline pressure. As such we draw a negative relationship from the former to the latter.

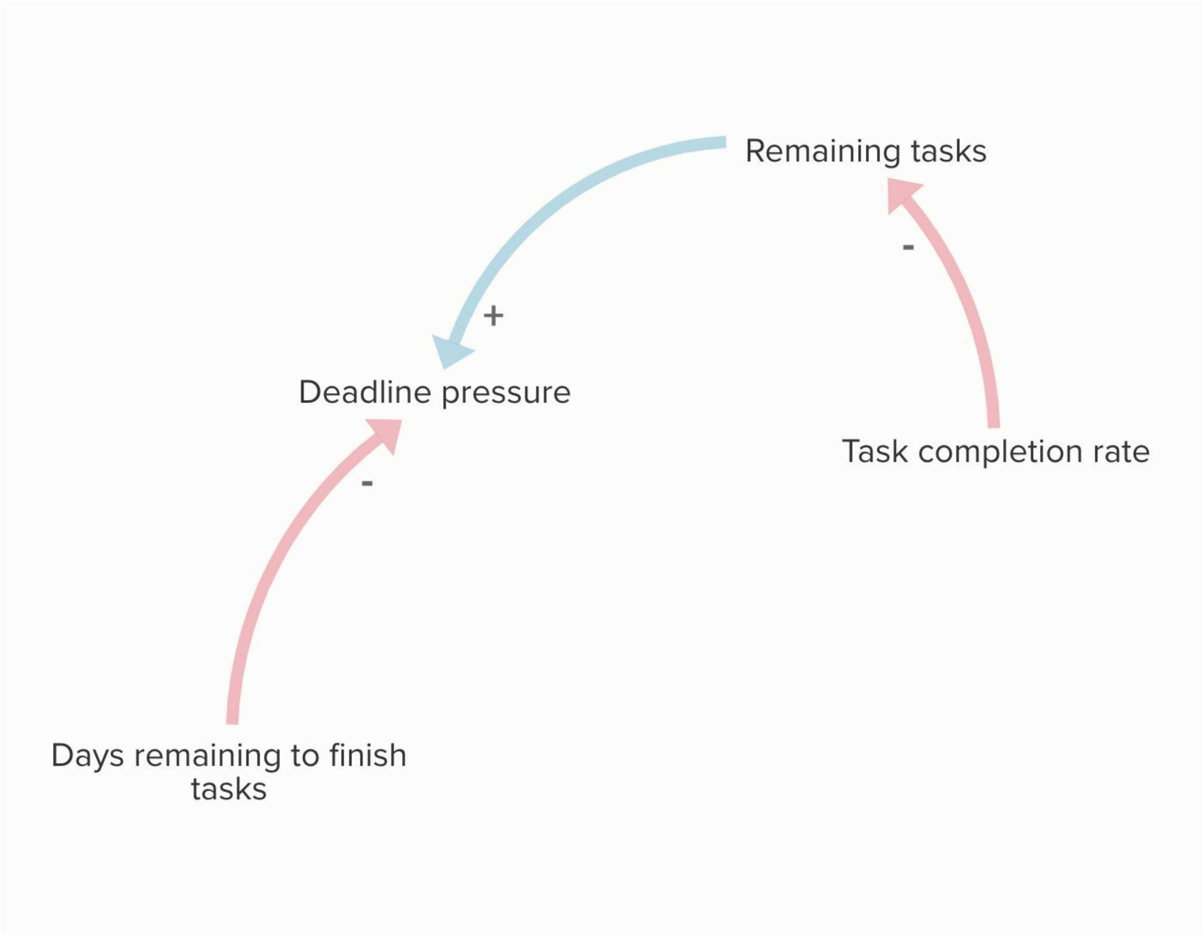

Exercise 4E: Completing the Model

What else remains for us to complete the model? Can you finish it yourself?

Answer to 4E

Mary spoke about how Deadline pressure increased her productivity, which we are calling the Task completion rate in our model. More Deadline pressure means a greater Task completion rate, and so there is a positive relationship going from the former to the latter.

Can you spot the feedback loop in this model and say whether it is balancing or reinforcing?

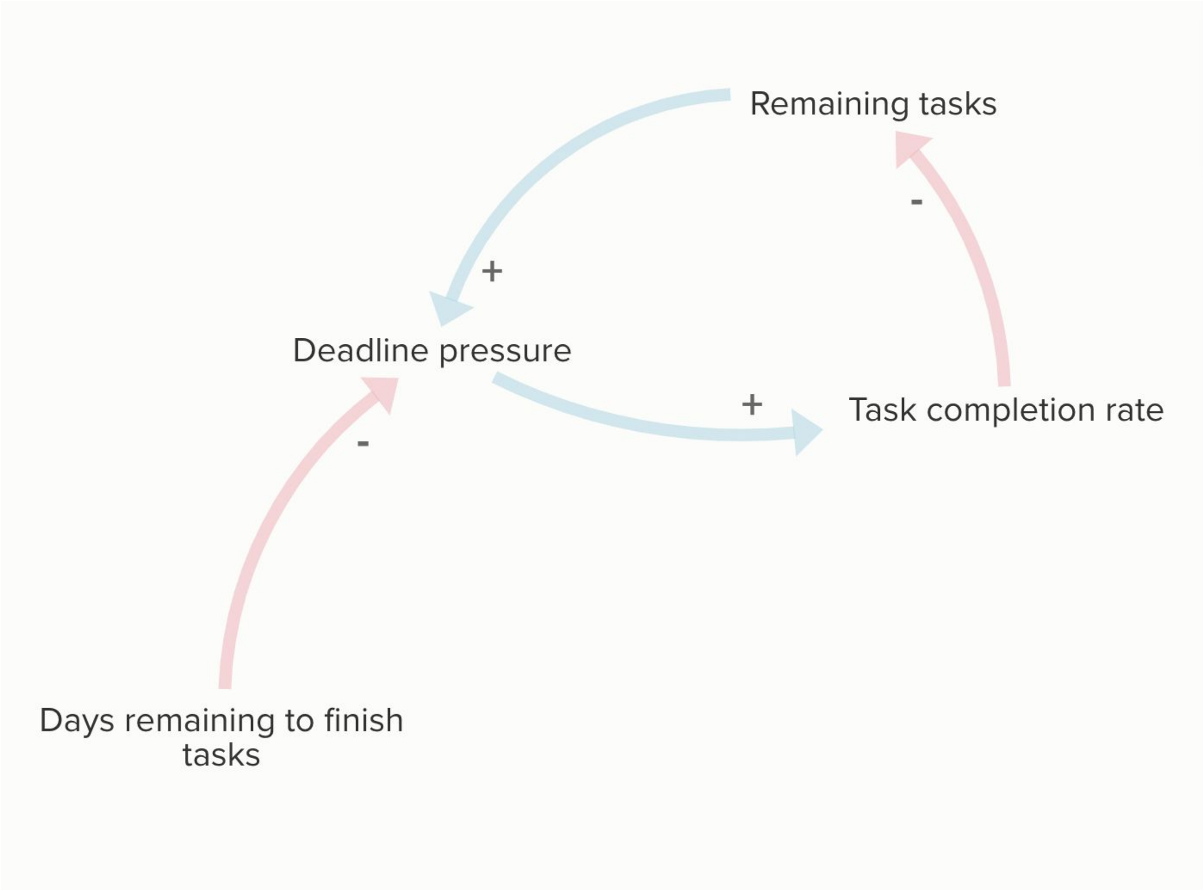

There is a balancing loop, which we are calling B1: Pressure means Productivity. This loop identifies how an increase in Deadline pressure results in an increase in the Task completion rate, which in turn leads to a reduction in Remaining tasks. This then reduces Deadline pressure.

Remember, to see all the variables in a loop, simply hover your cursor over the loop name (in this case, B1).

CLD Development

We so far have used 4 simple steps to develop our CLD diagram of Mary's mental model. To sum up, we have:

- Identified all variables that can be used to develop our CLD
- Chosen a central variable to start our model development
- Identified causal relationships to and from our central variable
- Completed our diagram by identifying all causal relationships to and from other variables whilst identifying important feedback loops.

In the next slides we will use the same steps to develop a CLD of Paul's mental model.

Exercise 4F: Paul's Mental Model

Let's try and use the steps on the previous page to create a CLD that represents Paul's mental model. As a reminder, Paul's answer to the question of how procrastination might affect productivity is shared again below:

Paul's Mental Model: Procrastination can increase errors and lower productivity

"I share Mary's perspective that when I'm under pressure I complete tasks quicker, but I would add that deadline pressure causes me to spend less time per task, which increases the amount of errors I make. More errors means more fixing errors, and so more remaining tasks to complete. I also sometimes miss these errors before submitting, which lowers the quality of my work. As such I believe I am better off starting an assignment early and completing it with less pressure and therefore less errors made."

Try to choose the variable names that you need to build Paul's mental model now. The variables we chose are on the next slide.

Remember:

- Choose only those variables that are necessary to communicate the essential message of the model
- Look for relationships and causality, and let them guide you to pick up important variables. Relationships will involve two or more variables
- Do not include verbs in your variable name; verbs are actions and they are represented by the causal arrows
- Use clear nouns to represent variable names
- Variables are usually referred to in the text multiple times with different names; try to group them into one variable name.

Answer to 4F: Identify Causal Relationships

Since Paul says that he agrees with Mary's point about how deadline pressure increases his productivity, we can keep the same structure as we had for Mary's CLD, and build on top of it. To the right you will see the new variable names that we have chosen to include in the model. If you wish to see why we chose these names, what viable alternatives there were, and what names would not be considered best practice, please click on the icon of each variable.

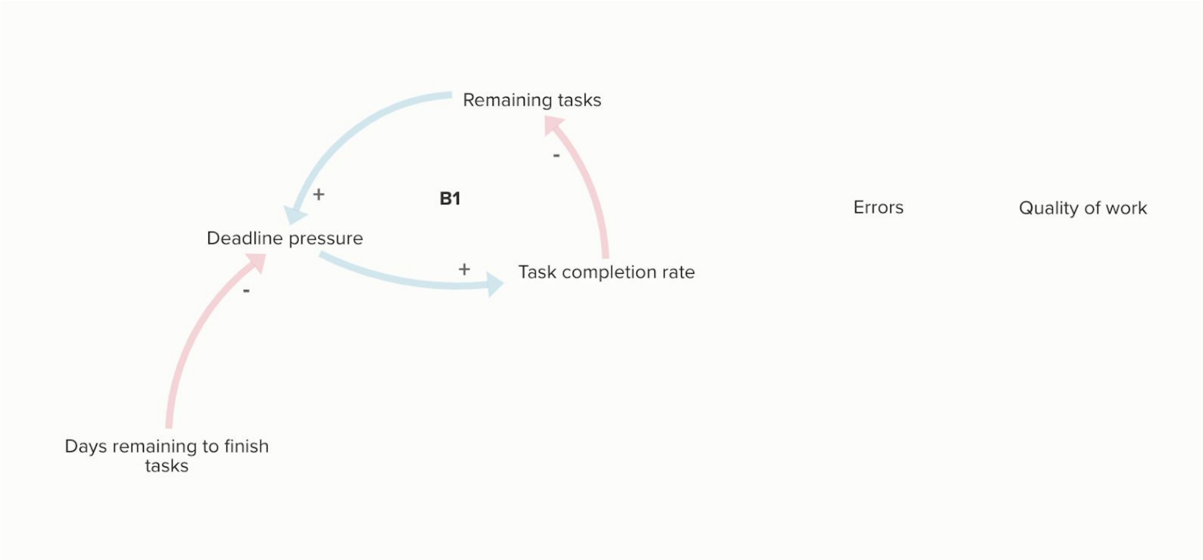

Exercise 4G

As a next step, try to draw all the causal connections that Paul mentioned.

Answer to 4G

Paul has said that when he completes tasks faster, he makes more mistakes. As such we can say that there is a positive relationship going from Task completion rate to Errors. At the same time, Paul has said that "more errors means more fixing errors, and so more remaining tasks to complete". We represent this by putting a positive relationship from Errors to Remaining tasks.

Paul has also said that Errors can lead to a lower Quality of work, given that he might not spot some errors before submitting the task. This is why we put a negative relationship going from Errors to Quality of work.

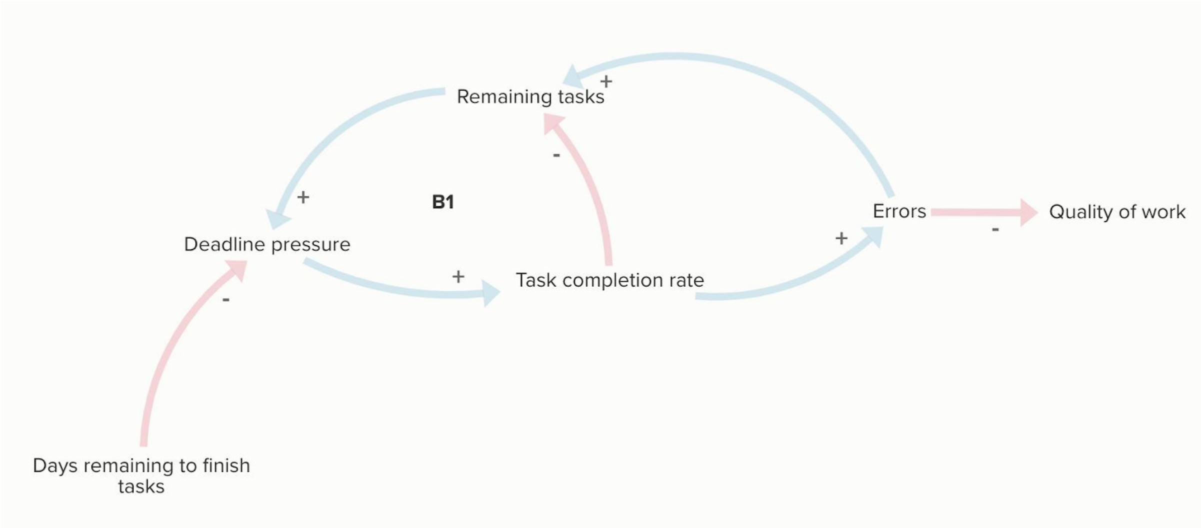

Exercise 4H

Can you spot the new feedback loop that exists in our model? And can you determine whether it is a reinforcing or a balancing feedback loop?

Answer to 4H

Below we see that there is a reinforcing feedback loop (remember, to see all of the connections in this loop, simply hover your cursor over the R1 icon).

This is an important loop to note. Whilst B1 shows the mechanism by which procrastination can increase productivity, R1 shows how it can reduce it. R1 shows how waiting for the Deadline pressure to kick in can increase your Task completion rate, such that you make more Errors. This can not only potentially hamper the Quality of work, it can also increase the Remaining tasks you have, which further increases Deadline pressure and so further increases the Errors you make! This is a vicious cycle that is good to be aware of. Paul thinks it is true for him. However maybe Mary thinks differently; maybe she doesn't make errors under pressure!

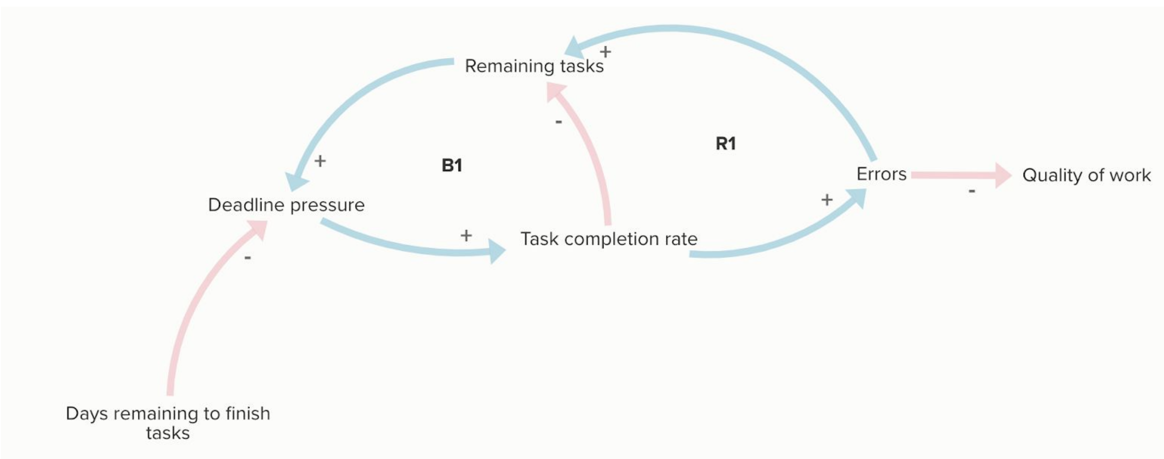

Insight: Feedback loops help us analyse the underlying structure of a system that creates the dynamic behaviour we observe.

End of Section 1: Modelling Procrastination and Productivity

The point of the previous section was to guide you through the process of building a CLD. We thought it was good to start with a model that is somewhat personal and perhaps easy to relate to. In the next section, we are going to zoom out a bit and look at an issue that is more macro in nature.

Section 2: Modelling the Utility Death Spiral

It's generally agreed that the shift to renewable energy is essential if we are to combat global warming. However, we should still be mindful of how we go about that shift, and any potential unintended consequences we create along the way. Below is an excerpt from an article describing the current situation that electric utility companies (companies that produce and distribute electricity) are facing as a result of the increasing adoption of distributed generation* technologies such as rooftop solar panels:

"It's a real-world example of the 'death spiral' that the industry has so far only considered in theory: as grid maintenance costs go up and the capital cost of renewable energy moves down, more customers will be encouraged to leave the grid. In turn, that pushes grid costs even higher for the remainder of customers, who then have even more incentive to become self-sufficient."

Before going deeper into the issue by reading the article (linked a few pages ahead), we challenge you to develop a simple CLD that highlights the cause and effect relationships, and resulting feedback loop(s), explained in that excerpt.

To makes things clearer, we suggest that you focus your model on how the adoption of rooftop solar panels (rather than other distributed generation technologies) will affect the price of electricity for those still on the grid. Try to sketch out your model now and then see how it compares to ours!

**Distributed generation refers to a variety of technologies that generate electricity at or near where it will be used, such as solar panels and combined heat and power.*

Exercise 4I: Identify the Variables

Try and identify some of the variables you could use to model the death spiral, as described in the excerpt (shared again below):

"It's a real-world example of the 'death spiral' that the industry has so far only considered in theory: as grid maintenance costs go up and the capital cost of renewable energy moves down, more customers will be encouraged to leave the grid. In turn, that pushes grid costs even higher for the remainder of customers, who then have even more incentive to become self-sufficient."

The variables we initially chose for this model can be seen on the next page.

Answer to 4l

The below is the initial set of variables we chose for our model. Please go to the next page to see how we initially drew the causal connections.

1 *Grid maintenance costs*

This variable represents how the cost of maintaining the grid changes over time

Viable alternative names: *Grid costs*

Incorrect names: *Increasing grid costs*. Variable names should have a neutral sense of direction, as the causal arrows allow us to represent how they might increase or decrease over time.

2 *Capital cost of solar panels*

A variable tracking the capital cost of solar panels

Viable alternatives: *Capital cost of renewable energy*. The excerpt mentions the capital cost of renewable energy. However we believe that the model will be less abstract and more easily understood if we focus on just rooftop solar panels, which are easy to imagine.

Incorrect names: *Capital costs* / *Costs*. Variables names should not be ambiguous. *Capital costs Costs* could mean the cost of the grid or the cost of solar panels. As such it's good practice to specify in the name.

3 *Customers on the grid*

A variable tracking how many customers are on the grid at any given time

Viable alternative names: *Grid customers*

Incorrect names: *Customers* / *Remainder of customers*. Again, variable names should be explicit rather than ambiguous.

Exercise 4J: Identify Causal Relationships

Before we go further, can you draw the causal connections between these variables yourself? Remember to base your model on the below theory:

"...as grid maintenance costs go up and the capital cost of renewable energy moves down, more customers will be encouraged to leave the grid. In turn, that pushes grid costs even higher for the remainder of customers, who then have even more incentive to become self-sufficient."

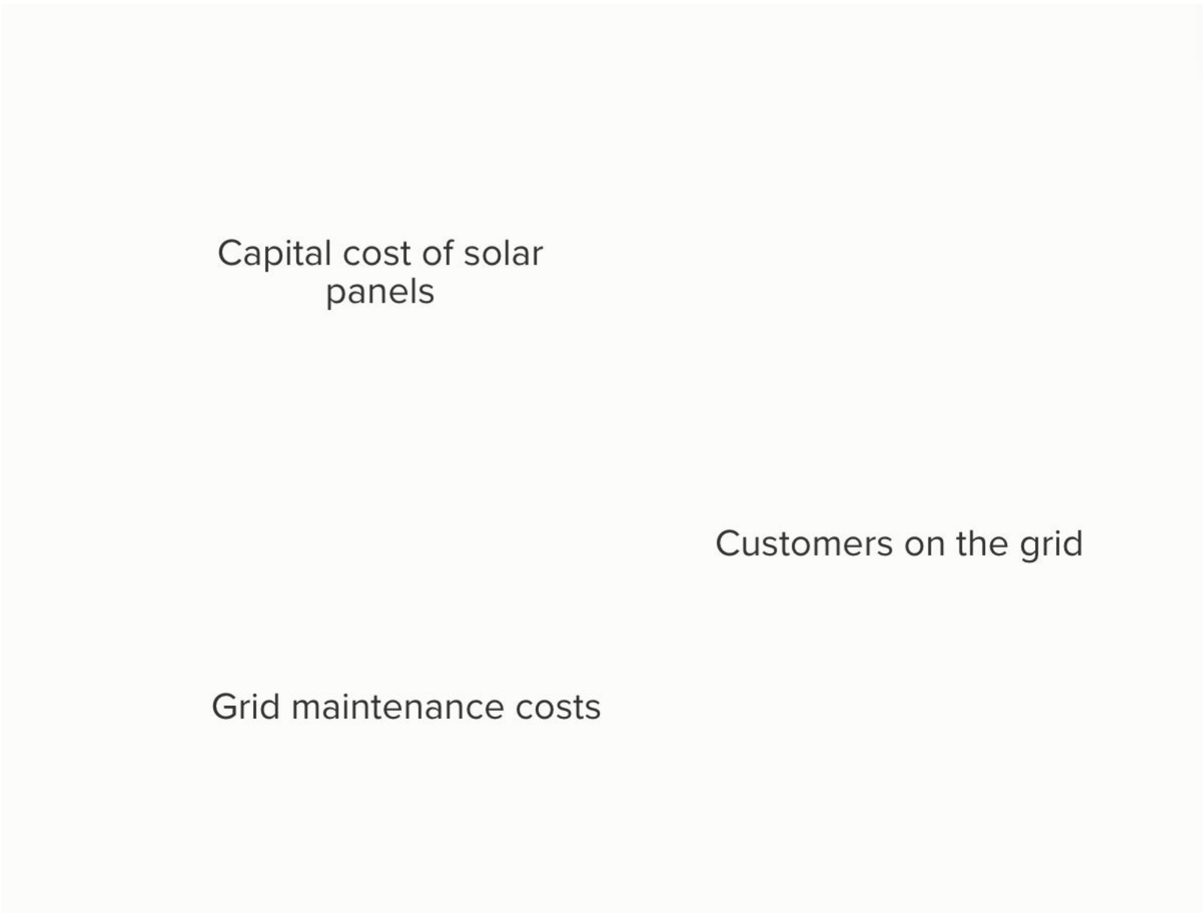

Answer to 4J (1/3)

"As grid maintenance costs go up... more customers will be encouraged to leave the grid."

As such there is a negative relationship going from Grid maintenance costs to Customers on the grid. The theory here is that higher costs means higher prices, which means that more customers will find it economically attractive to leave the grid and generate their own electricity.

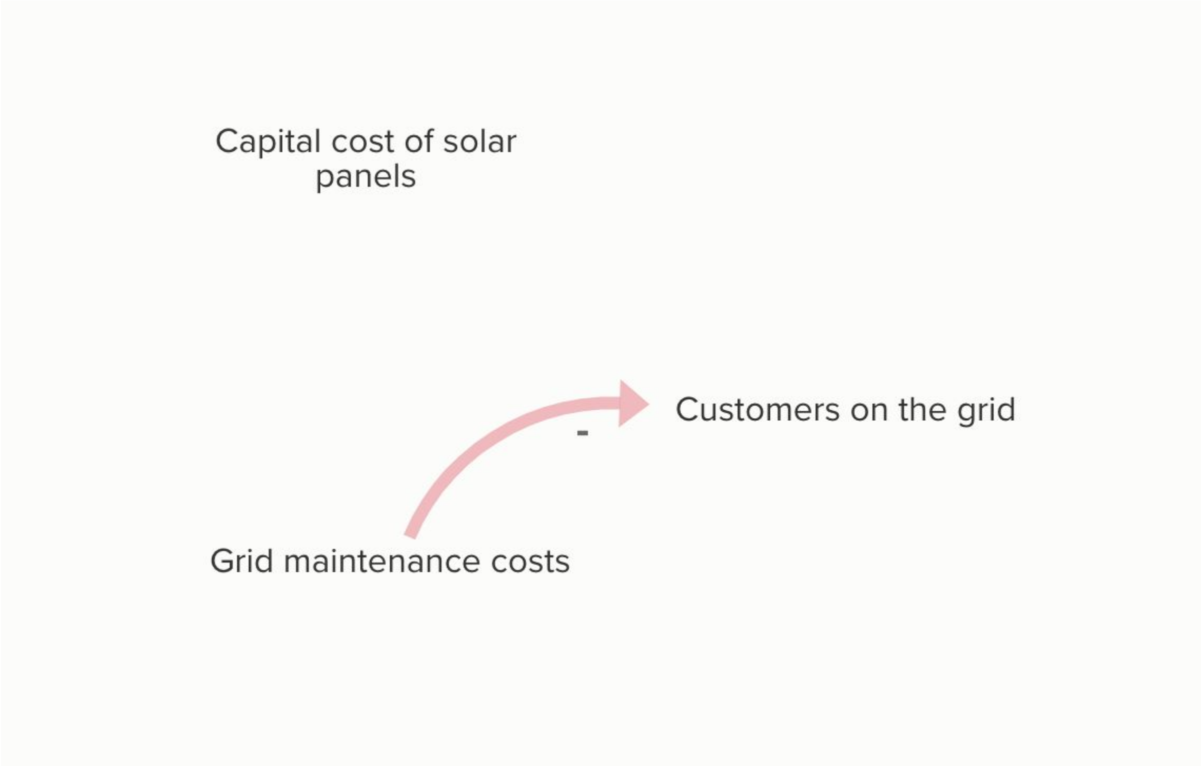

Answer to 4J (2/3)

The excerpt also mentioned that

"as ... the capital cost of renewable energy moves down, more customers will be encouraged to leave the grid."

This means that as the Cost of solar panels goes down, then there will be fewer Customers on the grid, because the cost of self-generation has gone down, meaning again that it is more economically attractive to generate your own electricity (through rooftop solar panels, for example).

Note that you could have another variable here called Customers leaving the grid, which would have a negative relationship to Customers on the grid. However we chose to leave this variable out for simplicity.

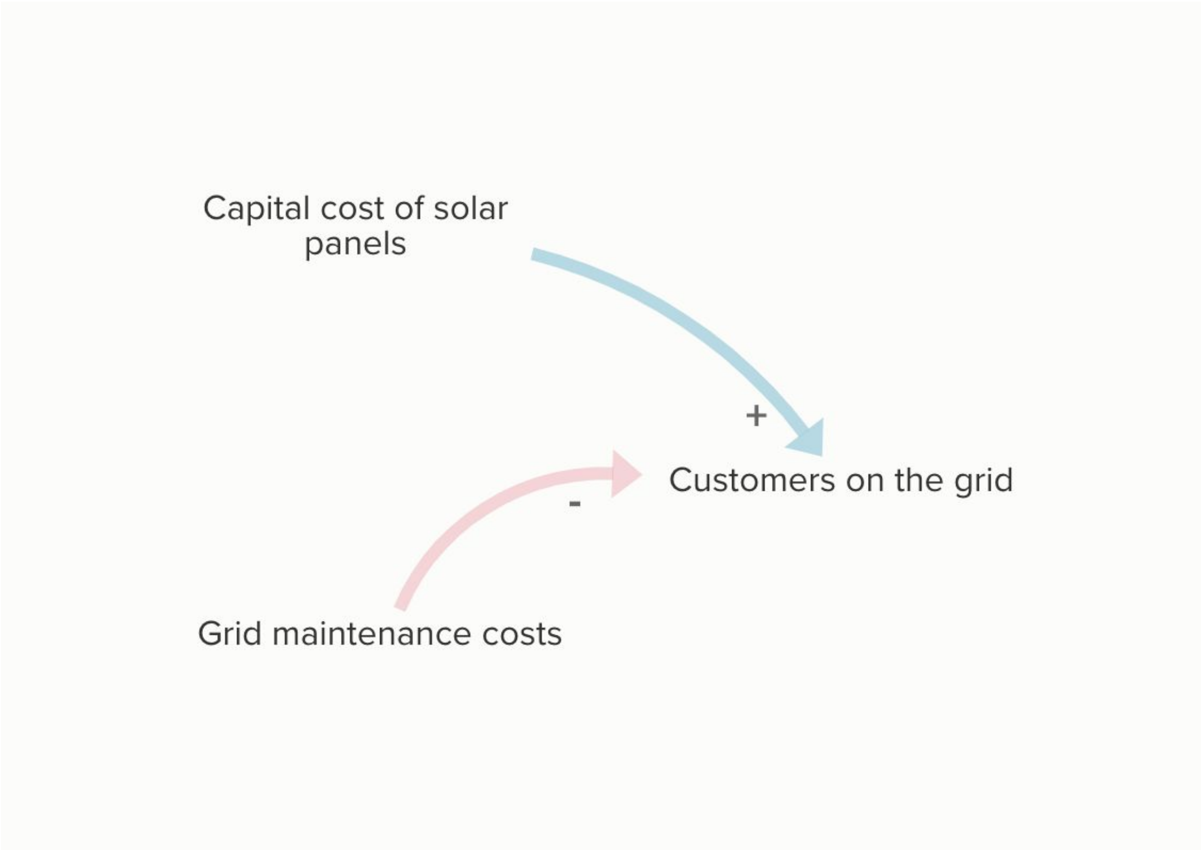

Answer to 4J (3/3)

The excerpt has said that as customers leave the grid, "... that pushes grid costs even higher for the remainder of customers, who then have even more incentive to become self-sufficient."

As such we drew a negative relationship going from Customers on the grid to Grid maintenance costs. This is because maintenance costs are relatively fixed costs, that aren't proportional to the number of people using electricity from the grid. As such, less people on the grid means that the utility company needs to recover its fixed maintenance costs from fewer and fewer customers, which they do by increasing prices. This in turn causes more customers to leave the grid.

This creates a reinforcing loop in the system, which explains why the problem is known as a death spiral!

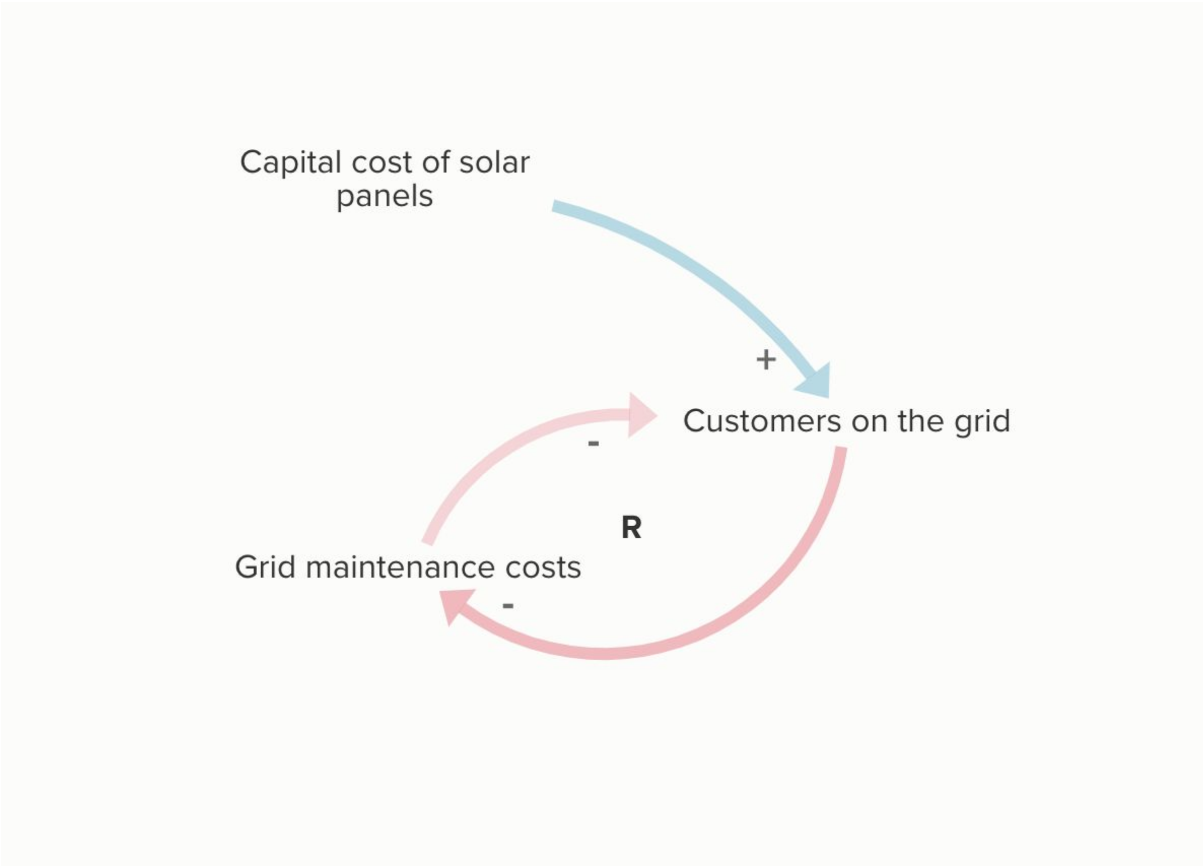

Critically Assessing the Model

Developing a CLD is always an iterative process, and we constantly need to reflect on whether our model is accurately representing what we/someone else thinks. Thus it's important to critically examine each and every causal link, to ensure that it is accurate.

Upon examining this model, we have noticed an error: a way in which it does not properly represent the mental model in the excerpt. Can you spot this error?

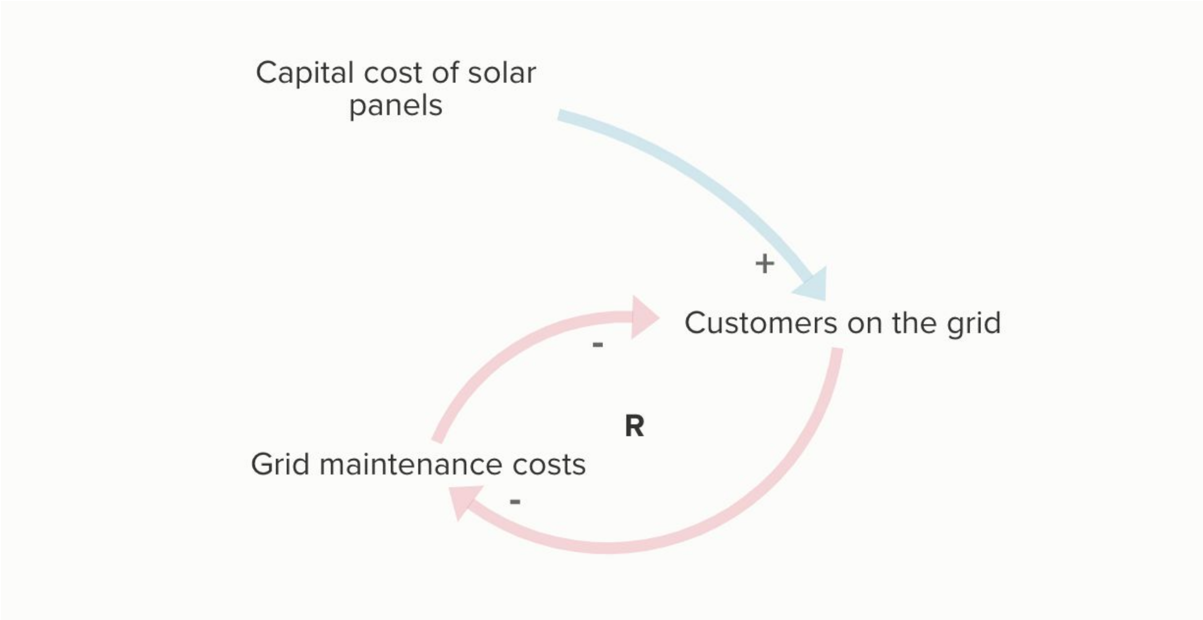

Exercise 4K

The excerpt has said that as customers leave the grid, "that pushes grid costs even higher for the remainder of customers.."

Is this accurately represented by the relationship on the right? We believe it is not. How would you best represent the relationships in a mode? Go to the next page to see how we chose to model it!

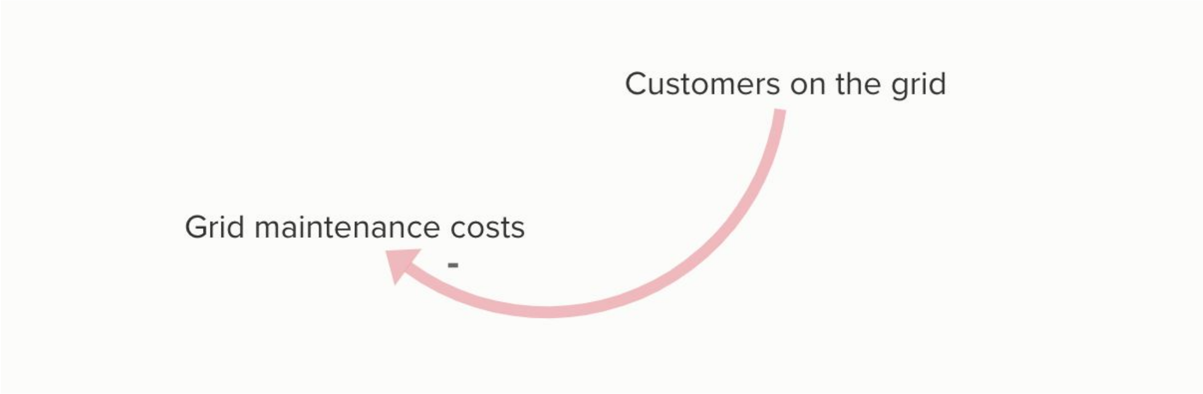

Answer to 4K

Given that the author mentions the costs being greater for the remainder of customers on the grid, we chose to model the situation as shown here. A decrease in the Customers on the grid will not decrease the Grid maintenance costs (indeed, it might even reduce it slightly). Rather, the author is saying that it will affect the Grid costs per customer, as less customers means less people to split the costs with.

Grid maintenance costs can then be modelled as another factor affecting Grid costs per customer. The higher the maintenance costs, the higher the costs per customer, so there is a positive relationship going from the former to the latter.

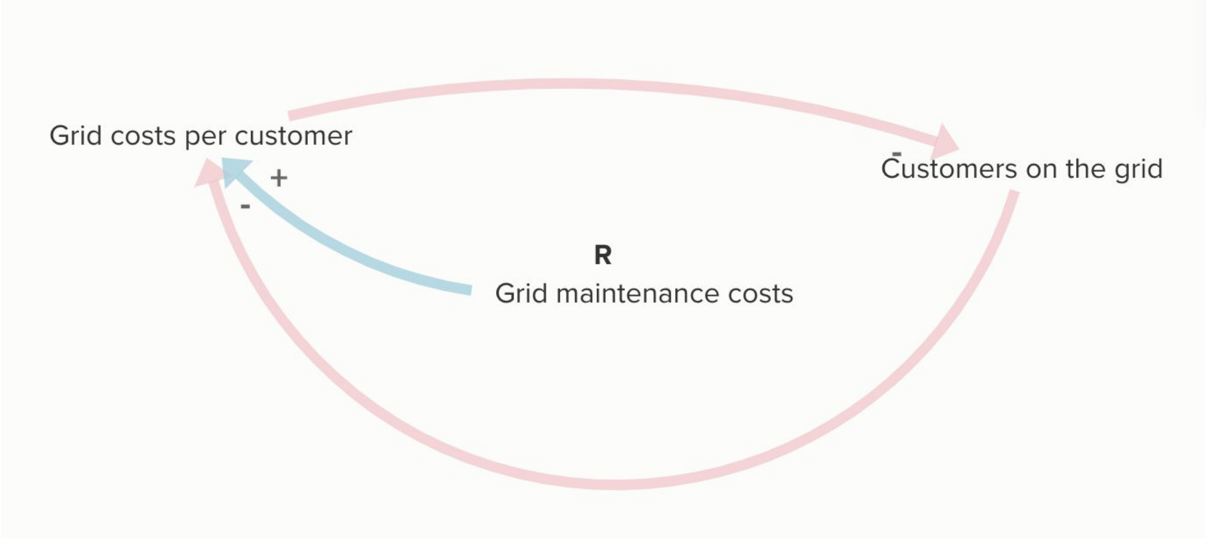

As such, now our model looks like this.

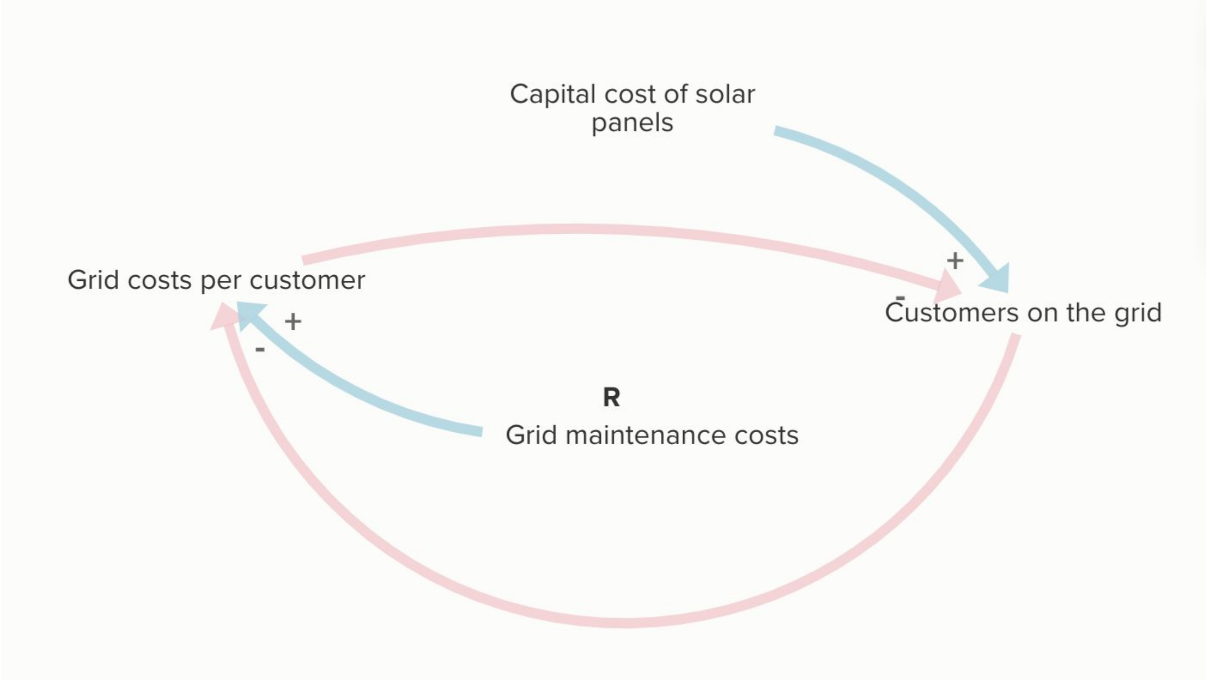

However, upon further discussion, we decided to make some more changes to the model, to make it more easily understood and more valuable.

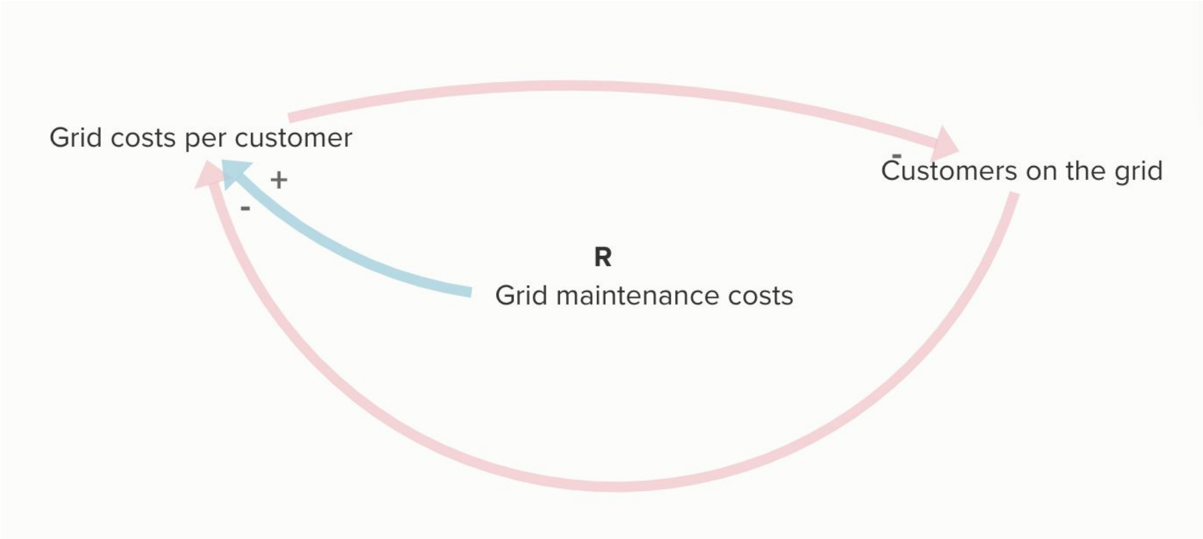

The article seems to assume that people leave the grid because there is a financial incentive to do so. As such, we decided to add a variable called Incentive to leave the grid because we felt it made this assumption more explicit and easily understood in the model.

The greater the Grid costs per customer, the greater the customer's Incentive to leave the grid. At the same time, if the alternative energy source decreases in price (i.e. if Capital cost of solar panels decreases), then that will also create a greater Incentive to leave the grid.

Insight
CLD development is an iterative process where we always assess and critically challenge our CLD so as to improve it.

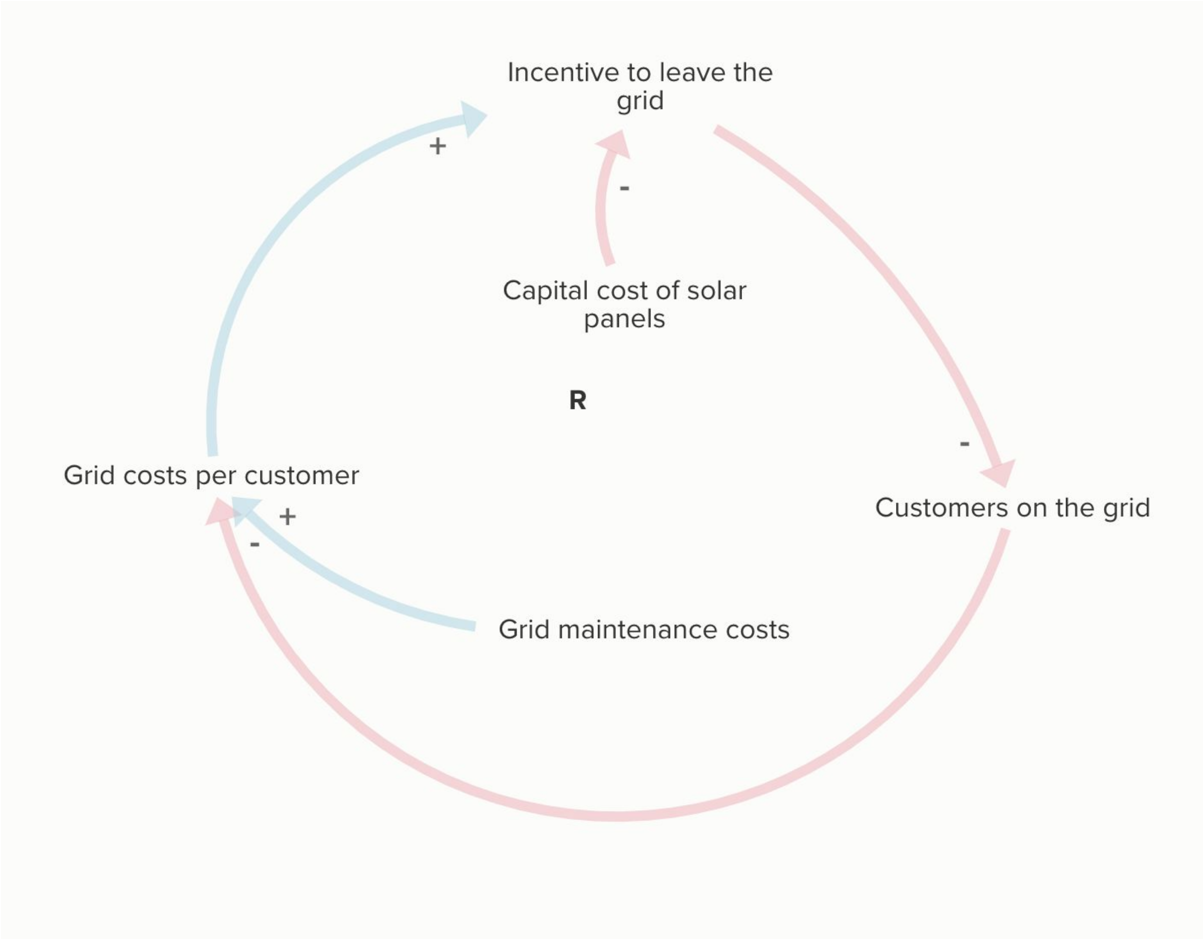

At this stage, we are satisfied that our model accurately represents the mental model expressed in the excerpt shared earlier.

However, CLDs are not only used to capture the meaning of what someone else has said. They are also great analytical tools. And as mentioned in Module Two, they can be great for highlighting assumptions in mental models, and allowing us to critically reflect on those assumptions.

In this case, we can use the model to discover and highlight the assumptions that are implicit in the author's mental model about the death spiral.

Before we show how we went about this, are there any assumptions that you notice yourself, which are implicit in the causal relationships shown in the model? And are there any changes you would make to the model in order to highlight those assumptions?

We noticed one assumption implicit in the negative relationship going from Incentive to leave the grid to Customers on the grid. This relationship assumes that households will be aware of this incentive, i.e., that they will be aware of not only increasing utility bills (which many people may not check), but also of the decreasing cost of solar panels. If households aren't aware of this incentive, then it will never cause them to leave the grid!

We chose to highlight this assumption by including an extra variable in our model, known as Awareness of incentive to leave the grid. As such, our model is now making clear that both the incentive and the awareness of the incentive will be important in influencing the numbers of customers on the grid.

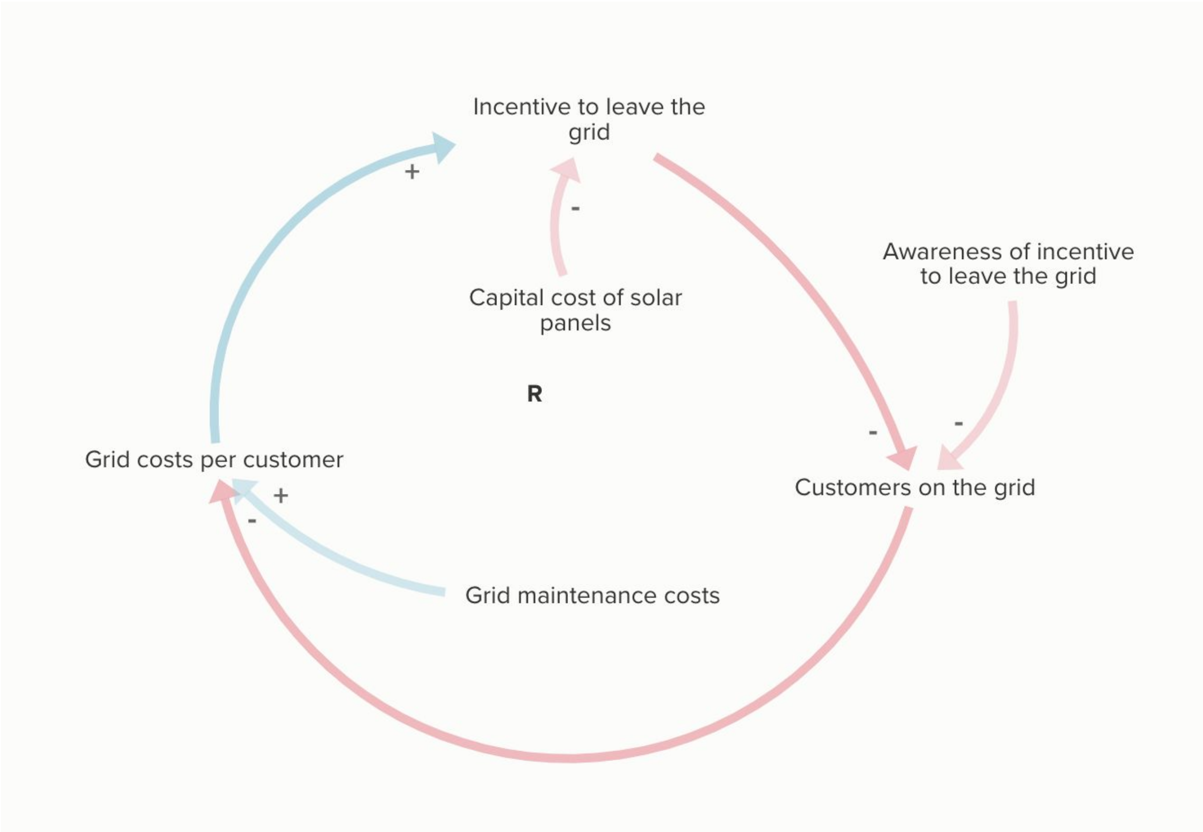

We also said that as the incentive increases, so should the awareness, as more and more solar panel salespeople are calling to your door, and are doing so with a better economic case for their product!

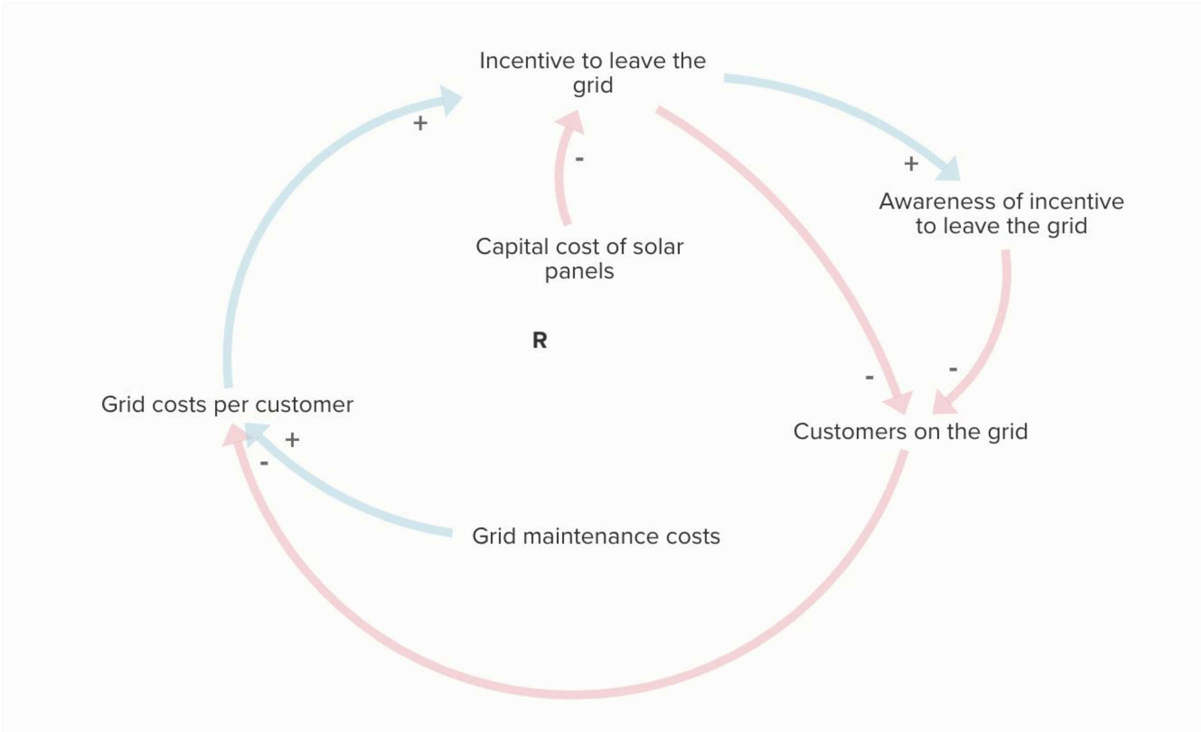

Another assumption we noticed in the model is that increasing Grid costs per customer leads to a higher Incentive to leave the grid. It seems that the author of the article is assuming that an increase in Grid costs per customer will lead to an increased incentive to leave the grid because utility companies will pass on increased costs to customers, in the form of increased electricity prices.

We chose to model this assumption by saying that an increase in Grid costs per customer leads to an increase in the Price of electricity from the grid. An increase in that in turn leads to an increase in the Incentive to leave the grid.

But do higher grid costs lead to higher prices? Now that we've highlighted the assumption, it's time to reflect on it.

We did some research to better inform ourselves. We found that, in most US states, the price of electricity is regulated. Each year, utilities go to a state body, known as the commissioner, and give them an estimate of the total costs they are likely to bear in supplying the demand for electricity in a given area. The commissioner then calculates the expected demand for a given area in the coming year, and uses that to determine the retail price that utilities will need to charge in order to achieve cost recovery.* So if demand falls due to more people generating their own electricity, and if costs do not reduce in proportion to lost revenues, then the commission will likely end up allowing the utility to charge a higher price for electricity.

**For example, if the expected costs are $10 million, and the expected demand for electricity is 10 million MWhs, then the commissioner will set the retail price of electricity from the grid at something like $1.01 per MWh (the extra $.01 per MWh allows for some profit for the utility* *and its investors). However if demand drops to 80 MWhs (due to people using alternative energy sources, for example), and the resulting expected costs of supplying this demand only drop to $90 million, then the commissioner will likely set the price as $1.125 per MWh (90 divided by 80, rather than 100 divided by 100).*

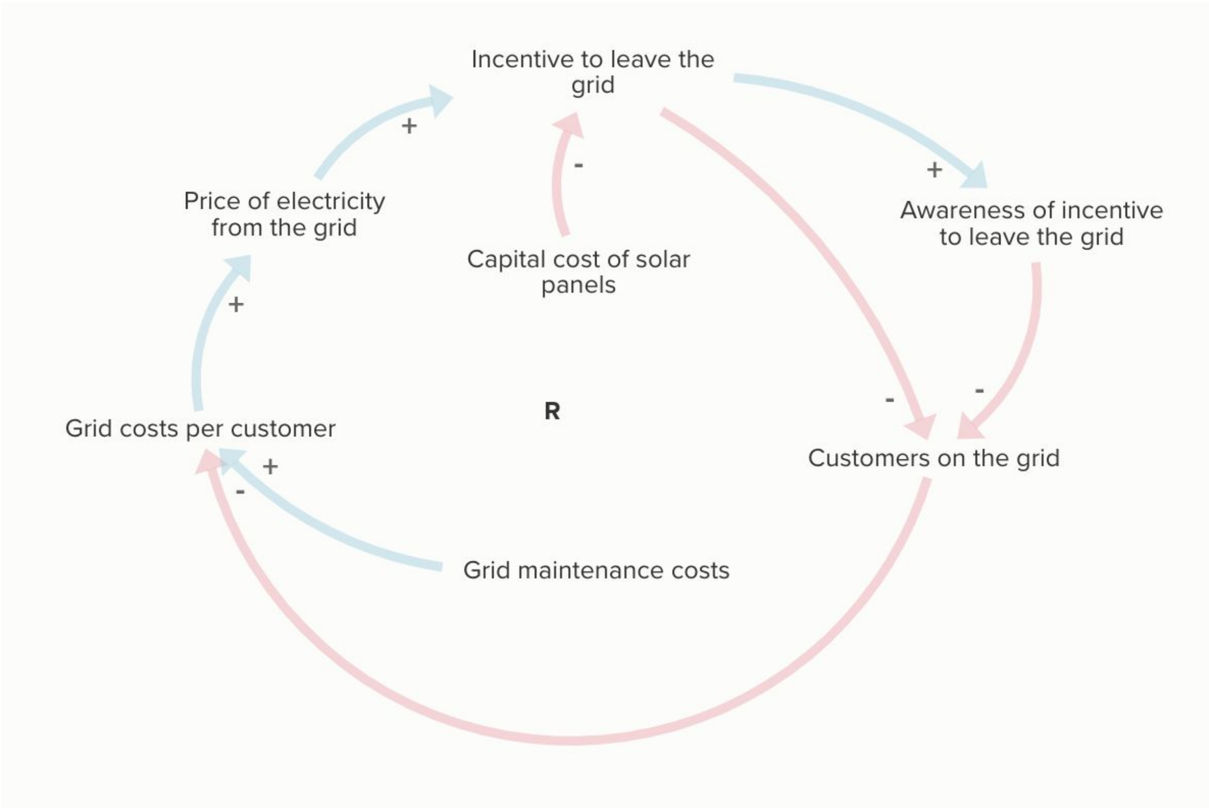

So now the question becomes: does a reduction in demand for electricity from the grid (due to rooftop solar use) decrease a utility's costs (through reduced production costs, for example) such that costs avoided compensate sufficiently for revenues lost?

Some articles we found suggest that if demand for electricity from the grid reduces—due to rooftop solar use, for example—then the utility's costs do not reduce in proportion to its lost revenues, because many of the utility's costs are fixed, such as maintenance of the grid.

However other articles we read suggested that utilities are greatly underestimating the benefits they experience from rooftop solar (and other distributed generation technologies). Such articles argue that rooftop solar brings many benefits that are not accounted for (for example, less stress on the grid, as electricity is consumed at the point of production rather than having to travel via the grid).

As such, they argue that rooftop solar adoption will decrease the utilities' costs in a way that is roughly proportional to the revenues that they lose as a result of rooftop solar use. This means that the increasing spread of rooftop solar will not necessitate the increase in the price of electricity from the grid.

These are all interesting points to note, but our aim here is not to provide an answer. Rather, we hope to have demonstrated that modelling this problem helped us to ask the right questions; it helped us to spot the assumptions implicit in the author's mental model and to critically reflect on those so as to better inform ourselves about the validity of a certain perspective (in this case, the death spiral hypothesis).

In cases such as these, CLDs can help us to get a better grasp of a problem, and then communicate our increased understanding (and our assumptions) to others. This CLD could also be used to analyse certain policies to address the issue of the death spiral, although that might require expansion of the model, and would likely require us to develop a quantitative simulation model.

A Note on Determining what is Included in Your Model

As mentioned in Module Two, models are abstractions of reality. As such, they cannot represent all the details of reality. This means that you must decide on what to include and what not to include in your model.

There are essentially two things to consider here: the boundary and the resolution of your model. We will discuss each of these now.

A Model's Boundary

The boundary of a model means the topics/issues that you will be focusing on in your model. For example, in the previous model on the death spiral, within our boundary was the price of electricity from the grid, the cost of renewable energy, the number of people on the gird, etc. Other potentially relevant factors such as the price of fossil fuels were not included. We say that such variables were outside the boundary of the model. In this case, we considered these issues to be outside the boundary of our model because although they have an effect on the speed at which the death spiral may happen, they are not part of the death spiral hypothesis itself.

Deciding on your model's boundary is always done on pragmatic grounds – it all depends on what is useful for your model’s purpose. Not including important issues has obvious downfalls, but including too many simply clutters the diagram and takes attention away from the most important causal relationships you want to display. In the case of the death spiral model we decided that including more variables/topics (such as the price of fossil fuels) was not necessary and would probably have led to a CLD that would be too large to be useful, especially given that this is an introductory course to CLDs. Indeed, knowing your audience is another important factor to consider when deciding on a model boundary.

A Model's Resolution

Once your boundary is decided, you need to decide on the Resolution of your model, which means the level of detail with which you will represent certain topics/issues. A model with a higher resolution will show more of the causal mechanisms that link variables together, whereas a model with lower resolution will not give as much detail about the causal mechanisms in the system at hand.

For example, in the death spiral model, we would say that we increased the resolution for the model when we included the Price of electricity from the grid variable in between the Grid costs per customer and the Incentive to leave the grid. Showing the assumed causal mechanisms in greater detail here (i.e. increasing the resolution) helped us to better highlight and discuss the assumptions in the death spiral hypothesis. In this case, it highlighted how we assumed that higher Grid costs per customer would lead to a higher Price of electricity from the grid, which would then increase the Incentive to leave the grid.

If, on the other hand, you felt an increase in the Grid costs per customer would obviously lead to a greater Incentive to leave the grid, then perhaps you wouldn't need to include the Price of electricity variable. However we felt that this relationship was not so obvious and so was worth highlighting,

Lower resolution (assumptions less explicit):

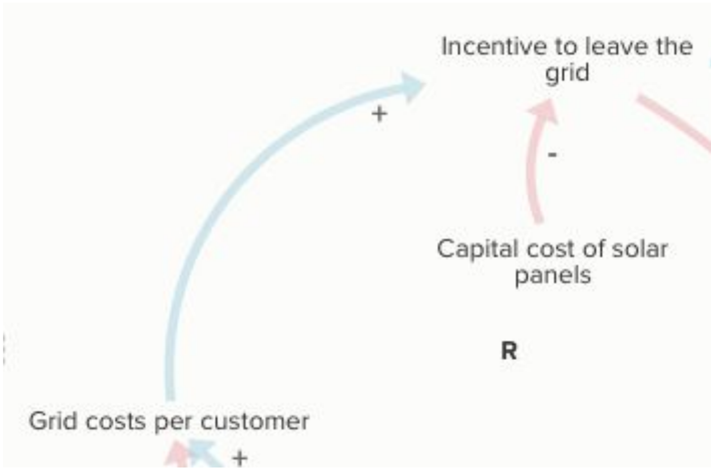

Higher resolution (assumptions are more explicitly represented):

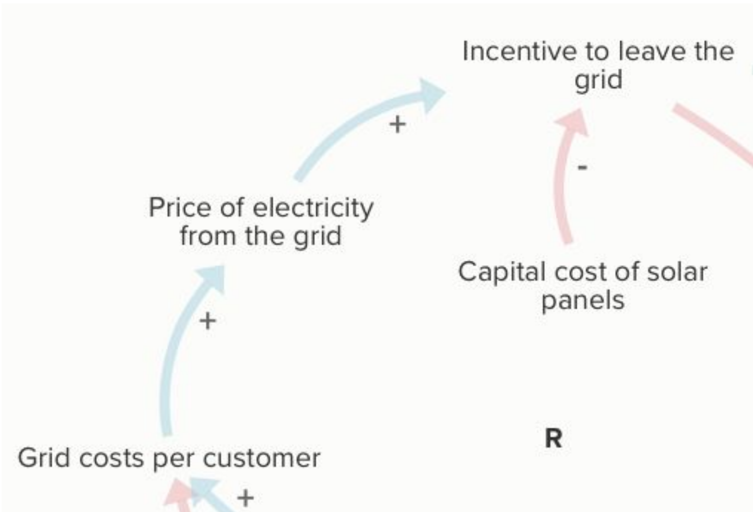

Again, deciding on the resolution of your model is done on pragmatic grounds: we want to avoid models that have too many variables, but we also want to include sufficient detail for representing important assumptions and causal mechanisms.

Determining a Model's Boundary and Resolution

Ultimately, there are no hard rules for determining the boundary and resolution that you set for your model. It is a decision to be made on pragmatic grounds, and it will depend on many factors, such as the reason for which you are making the CLD, and the audience you are creating it for.

Also, it's important to remember that determining the model's boundary and resolution is an iterative process, and we encourage you to consistently reflect on whether your CLD's boundary is too narrow/broad, and its resolution too high/low.

Keep this in mind as we move to the next section.

End of Section 2: Modelling the Utility Death Spiral

That brings us to the end of section two. We hope that by now you have learned a bit more about how to create a CLD that accurately represents someone's mental model/hypothesis, and that you are aware of some of the important considerations to be made when doing so. We discussed the importance of deciding on a boundary and level of resolution for your model, as well as the importance of constantly reflecting on your model and iteratively improving it. We hope to have also demonstrated how building a CLD of someone's hypothesis can help one to critically reflect on that hypothesis in a deeper way than we usually do.

Section 3: Modelling the Theory of the Invisible Hand

In this section, we are going to make things a little more challenging. Rather than provide you with text, we are asking you to watch a video and build a CLD of the theory it describes. The theory is Adam Smith's beliefs about how the price of commodities can be kept in balance when economic actors are allowed to act in their own self-interest. This theory has become known as the Theory of the Invisible Hand.

Please click on the video below to see an explanation of Smith's theory.

Exercise 4L: Modelling the Theory

We challenge you to build a CLD that succinctly explains the cause and effect relationships assumed in Adam Smith's theory.

We will begin presenting the model we built on the next page. Our model describes the theory using the example in the video of how different bakers of scones regulate the price of scones whilst acting in their own self-interest. We decided to keep it simple and focus just on this aspect, leaving the supply chains for scones outside the boundary of our model.

Answer to 4L

We began our model with a variable representing the Price of scones from baker A. We then asked ourselves what affects this variable.

Smith assumes that when there is an increase in demand for something, sellers will increase the price of their product, all else equal. As such, we model a positive relationship going from Demand for scones from baker A to Price of scones from baker A.

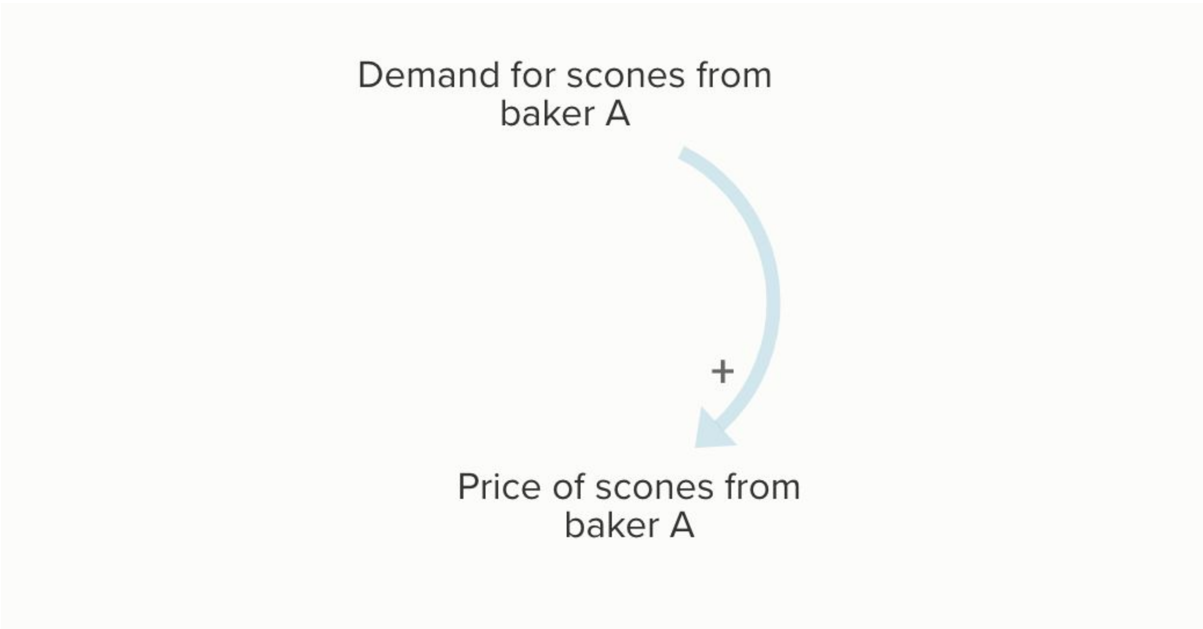

However demand alone will not determine the price of a commodity. According to the theory known as the Law of Supply and Demand, both supply and demand will impact the price of a good. The theory says that when demand for a product is high and supply is relatively low, then prices will be higher. The theory also suggests that when demand is low and supply is relatively high, the prices will be lower.

For example, if the baker knows he will have 20 people that want to buy his scones, but he only has 15 left, he or she might increase the price of the scones. Alternatively, if the baker makes too many scones in a given day, and has many left in the shop, then he or she might consider lowering the prices of scones (offering a two-for-one promotion, for example).

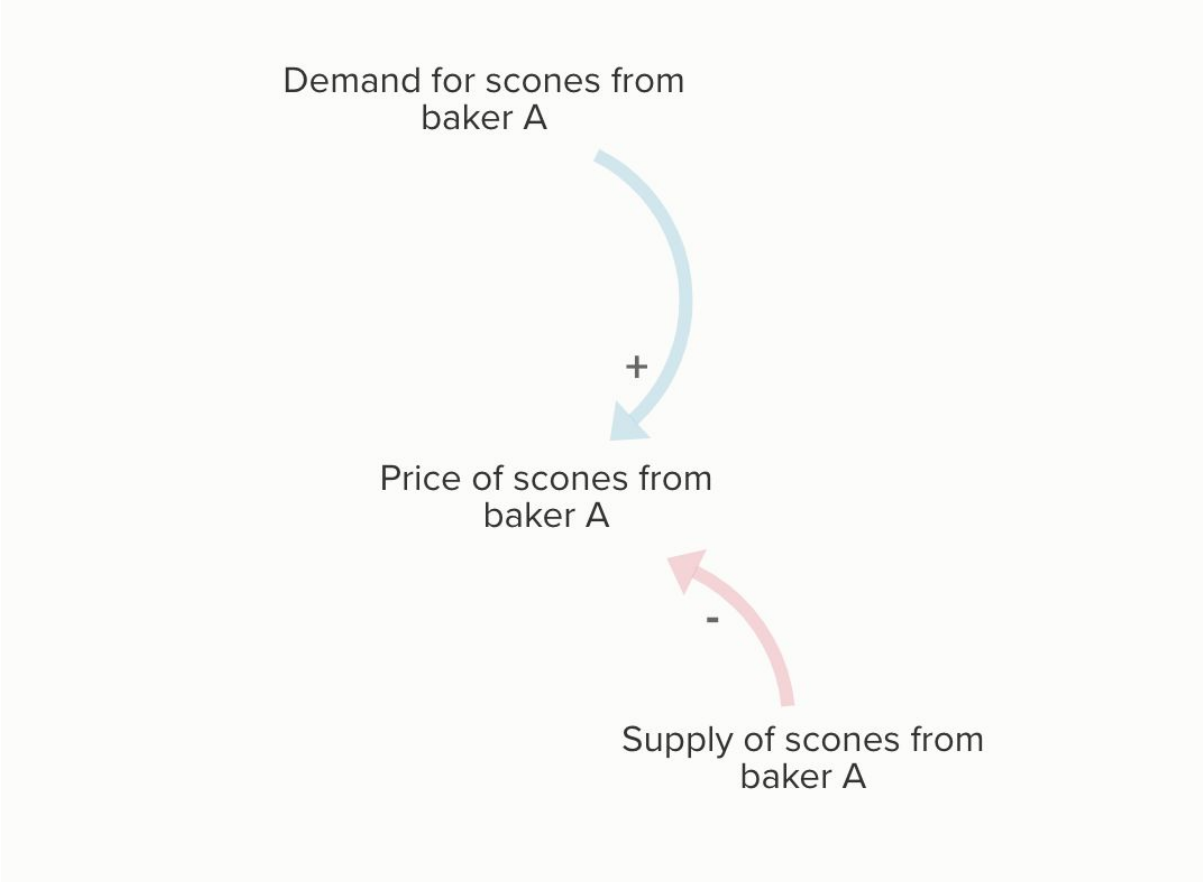

Next we can ask ourselves what determines the supply of scones in the first place. Smith's theory was about people acting in their own self-interest. As such, we can assume that the baker will only supply scones if there is a profit to be made in doing so. The greater the Potential profit for baker A, the greater the Supply of scones from baker A, as shown in our model.

The potential profit will increase when the price is higher, all else equal, and so we included a positive relationship from Price of scones from baker A to Potential profit for baker A.

We also decided to include a variable called the Cost of making one scone, which can be thought of as the cost of ingredients, labor, and all other factors needed to make one scone. The higher the cost, the lower the potential profit, all else equal.

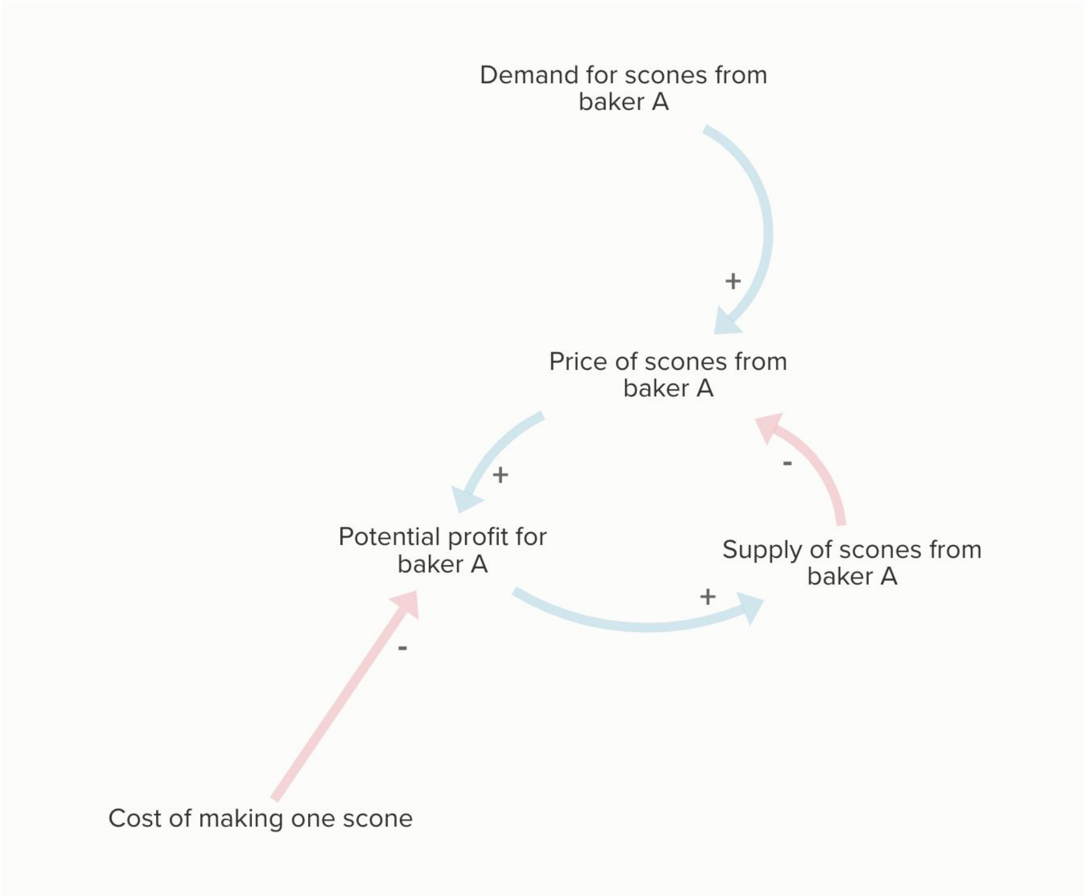

As such, we see the first balancing loop that explains part of the dynamics of Smith's theory. This loop says that the more profit there is to be made from supplying a good, the more of it will be supplied. And the more of it supplied, the lower the price, and so the lower the potential profit, and the lower the supply.

This loop helps to regulate the price of scones, and brings supply in line with demand.

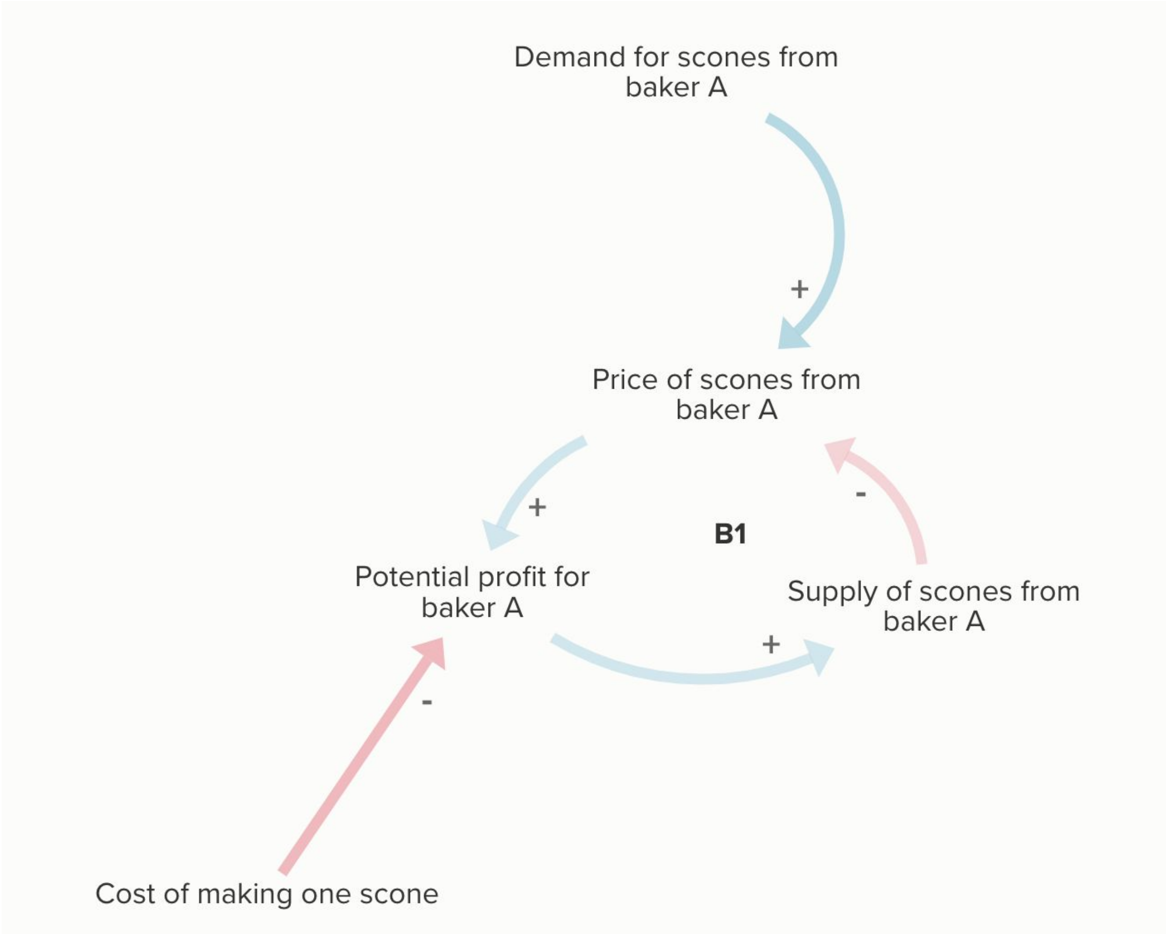

However, the theory doesn't end there. Smith also talks about how competition between different actors also helps regulate the price of goods. In the video, we saw how another baker noticed the potential profit to be made in supplying scones. As such, the dynamics for baker B can be modelled in the same way, as shown.

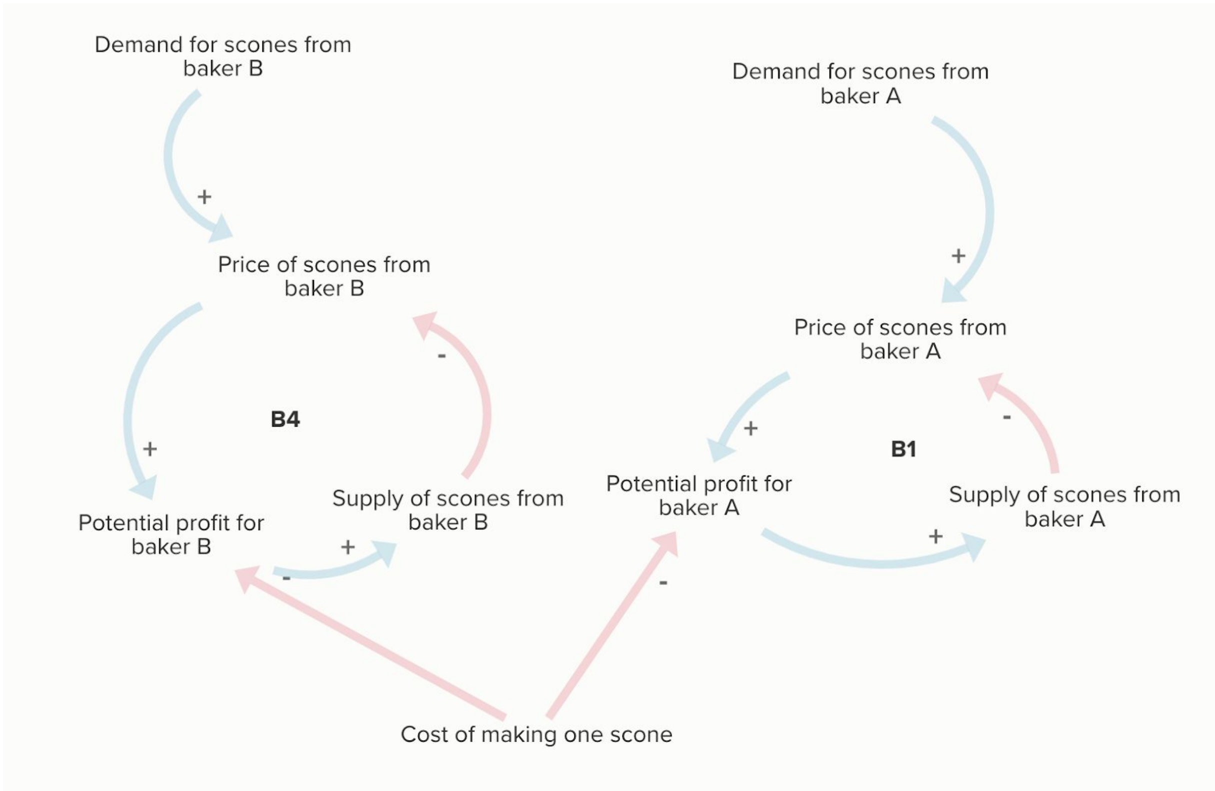

Now that there are two bakers supplying scones, there is competition in the market. This means that customers have a choice as to which baker to go to for scones.

To model this, we said that the prices of scones from each baker will determine the Relative value of scones from baker A: a higher price from baker A leads to a lower relative value, whereas a higher price from baker B leads to a higher Relative value of scones from baker A. Given that price is the only factor affecting the relative value of scones from each baker, we are of course implicitly assuming (for simplification) that the quality of scones from each baker is the same.

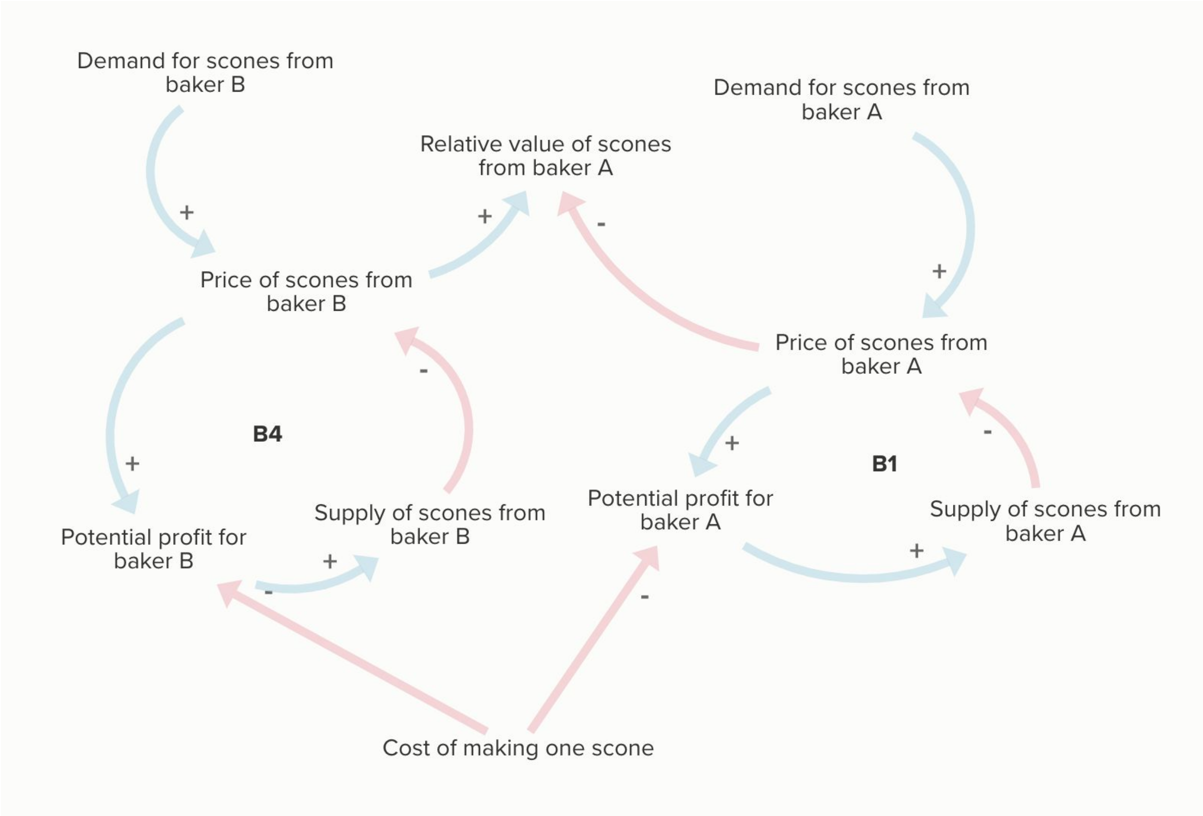

The Relative value of scones from baker A will then determine the demand for scones from each baker, as customers choose what they believe to be the best value. Note that this relationship in itself assumes that customers are aware of the price of scones from both bakers, which is not the case in all markets!

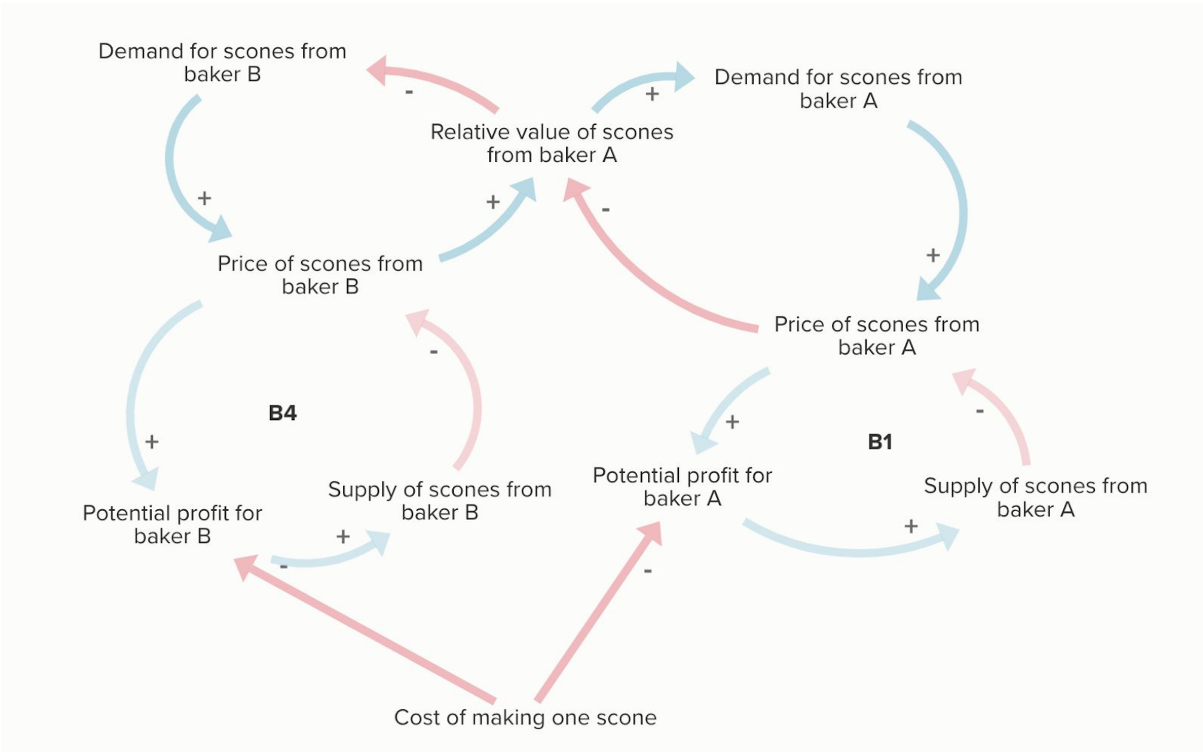

Now we see an additional balancing loop for each baker. Loop B2 says that as the Relative value of scones from baker A increases, so will the Demand for scones from baker A. This will in turn increase the Price of scones from baker A, which reduces the Relative value of scones from baker A.

At the same time, an increase in the Relative value of scones from baker A will cause baker B's demand to drop, which causes their price to drop, which in turn decreases the Relative value of scones from baker A, or in other words increases their relative value. As such, loops B2 and B3 describe the way in which competition between bakers regulates the prices that they each charge. When one baker increases their price, the other becomes more popular, causing the first baker to bring their price back down so as to remain competitive.

At the same time, neither baker is likely to lower the price of their scones so much that they no longer make a profit, as described by loops B4 and B1.

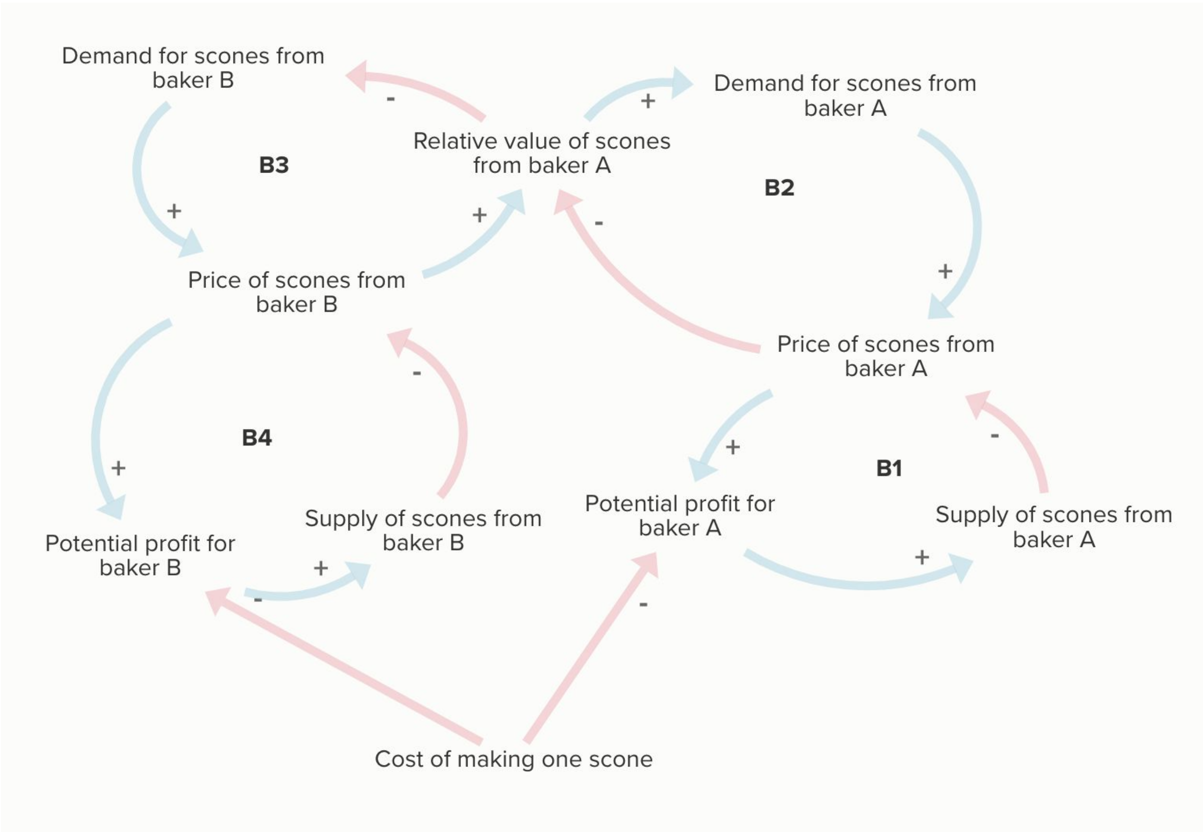

End of Section 3: Modelling the Theory of the Invisible Hand

We hope that in this section we have shown that CLDs can help us to better understand abstract theories by explicitly representing each of the cause and effect relationships that they assume. We appreciate that this was not an easy modelling exercise, but we felt that it was a good idea to give you a challenge!

Section 4: A First Experience of Group Model Building

In this section we will be giving you a challenge to complete with a group of people. This can be any group: your classmates, your colleagues, your friends, or even your family! The purpose of this exercise is to familiarise yourself with what is known as group model building (GMB), a methodology within system dynamics that we briefly introduced in Module Two.

As mentioned there, GMB involves engaging different stakeholders to a problem in a process of jointly building a model. The aim is to integrate different stakeholder perspectives, and allow for the development of a shared understanding of a problem. The process acts as an opportunity to improve communication between stakeholders, and the synthesis of different perspectives on a problem can often lead to the discovery of new solutions. Additionally, by involving stakeholders in the analysis, the group is more likely to be supportive of the outcomes of that analysis (i.e., the decision, policy, or strategy developed), which is extremely valuable in the implementation phase.

A Typical GMB Session

A typical GMB session involves identifying and gathering between 5-15 stakeholders in a room, together with at least two people experienced in system dynamics and facilitation. One of these system dynamicists will facilitate the group process, while the other will focus on building the model based on the group's conversation. The modeller can also interact with the group when necessary, for example to check that the model represents what the group has been saying. This allows the facilitator to keep her/his attention focused on the group dynamics, and allows the modeller to focus on producing an accurate and clear model that everyone can read.

The group usually starts by building a CLD of a problem, and may at a later stage create a simulation model if that seems necessary. It may take several half-day sessions until a group has constructed a model they can agree on, and from there it might take several more sessions to use the model for analysing potential solutions to the problem. In the end, the hope is that the group develops a shared understanding and a common motivation for implementing a certain solution/strategy.

Simulating a GMB Session

To get an idea of what running a GMB session is like, it is best to be involved in a real situation in which you and those around you are strongly invested in the outcome of the analysis. However, there is no harm in using a little imagination and running an exercise in which we imagine ourselves as stakeholders to a problem - which is what we will do in this section! We will use a real-life case of GMB that The Sustainability Lab were involved in, and will ask you and your group to imagine yourselves as different actors in this complex situation.

Before we explain more about what you will do in this exercise, it will be helpful to get a bit of background information.

The Martina Bustos Community

The case we will deal with occurred in 2014 and 2015, when fellows of The Lab's Global Sustainability Fellowship program were working together with the Martina Bustos community. This is a settlement that is located 4 km from downtown Liberia, the capital of the Guanacaste region in Costa Rica. The settlement was established about 15-20 years ago when a private landowner, named Martina Bustos, donated the territory so that low-income families could have a place to live. However, Martina never formally handed over the land titles, and to this day they remain with her children.

Nonetheless, many poor migrant families have moved to the area, and the population on the settlement has grown to between 3000-3500 people, mostly of Nicaraguan descent. The lack of land titles has caused many problems as, without them, the community is not formally recognised by the Costa Rican government. As a result, there is no legal obligation for the government to provide access to basic services such as electricity, waste management, and clean water.

In the summers of 2014 and 2015, participants in the Sustainability Lab's Global Sustainability Fellowship program began working with the Martina Bustos community. The purpose of the work was to create a community development plan that would serve as the blueprint for creating a more robust and sustainable community. The fellows worked with members of the community (mainly the women's association) to identify central problems as perceived by the community members, and try to find a strategy for resolving them. Much work had already been done by other NGOs in this area, but The Sustainability Lab's approach was unique in that it took a holistic/systems perspective on the community, addressing multiple problems at once (economic, social, environmental, and motivational) rather than focusing on just one issue at a time (for example, just housing or just water).

The Use of CLDs

As part of this work, the fellows held GMB sessions with the community members, the purpose of which was to discover together how the problems of the community might be interlinked. This analysis would then aid in the development of a holistic strategy for addressing the community's issues. For these sessions, about ten community members were invited to the local university. They participated in a two-hour session which was facilitated in Spanish by Vanessa Armendariz, one of the fellows.

Role-Playing

For this exercise, we want you and your group to split up into different roles, each of which is described below. If the group is larger than 12 people, then we suggest that you form two groups and conduct the exercise separately.

Role 1: Members of Martina Bustos (2 to 10 people): Your task is to put yourself in the shoes of members of the community. To become familiar with the nature of the issues faced by the community, we will provide you with some quotes from the community members (shared a few pages ahead). To keep things simple, you should try focus only on the issues in those quotes, so that the scope of the model does not become too broad to handle. Your job is essentially to experience a GMB session as if you were a participant, and we hope that this will be enlightening for when you become a facilitator or modeller.

Role 2: The Facilitator: You are essentially leading the group model building session. Your role is to facilitate the process from start to end, managing group dynamics and ensuring that all participants are able to express their views, and that the conversation maintains its focus on the issues at hand.

Role 3: The Modeller: Your role is to create a CLD based on what the members of the community are saying. While the facilitator will do most of the interacting with the group, you may interject every now and then when you need to clarify what someone has been saying, and whether your model is reflecting that accurately.

Logistics

It will help if you can do this exercise in a room with a whiteboard upon which you can draw the model, but if that is not possible then you can use a large piece of paper that the modeller can easily show to the other participants. We suggest that the facilitator stands in front of the group, and arranges the others in a semi-circle in front of her/him, as this creates the most balanced way for a group to interact. Ideally, the modeller would stand behind the facilitator, at the white board. Alternatively, if the modeller chooses to build the model on a computer then we suggest that they use a projector/digital display to make their model easily visible to the whole group.

Focus of the Session

To keep things simple in this exercise, we can imagine that the agenda for the meeting has been set around looking at issues related to the community's lack of land titles. We suggest that those playing the role of community members maintain their focus on just that issue, to avoid too many topics being discussed at once.

One of the facilitator's responsibilities is to keep the discussion focused, and if one particular issue has been agreed upon before the session, then that should help in determining what is off and on topic. Keeping the focus on an agreed boundary will allow for a more productive session, rather than letting the group's discussion touch upon many topics without delving into any in depth.

Of course, you and your group are unlikely to know much about the land titles issue of the Martina Bustos community. To help you play their role, you will find a number of quotes from participants in the actual GMB session on the next page. You can take inspiration from these quotes to try to put yourself in the shoes of the community members, and see the problem from their perspective. You can then relate these issues to the facilitator, who together with the modeller will try and develop a CLD that combines together the different perspectives.

We suggest that the people playing the role of facilitator and modeller do not look at these quotes, as part of the challenge of facilitating a GMB session is dealing with and understanding the issues expressed by participants as they arise, without any prior knowledge of them. So if you are playing the role of the facilitator or modeller, please go ahead and skip the next page and move to exercise 4M.

Quotes from the Community Members

Below is a series of quotes from the Martina Bustos community members during their GMB session with us in 2015. Please note that these have been translated from Spanish and simplified for educational purposes. For confidentiality purposes, we have not specified the author of each quote.

"I don't really feel like I belong in this community. To feel belonging, you need to feel a sense of stability in a place. Without the land titles, I feel that I could be kicked out of my home at any time. I think that affects how we see each other here: the lack of land titles means there is no feeling of community."

"At the moment, there is too much conflict in the community for progress to happen. Only when we work together will we achieve any real improvements."

"I think we need the landowner to hand over the land titles before the people here commit to improving the infrastructure in the community. Without the titles, there is no sense of community, and without the sense of community, there are no successful projects carried through."

"The landowner has said time and time again that he does not think we are ready to have the land titles. He says that until we act like a community, he will not hand over the titles."

"The government have said they will not support us with funding until we have demonstrated that we can work together on projects and see them through to the end."

Exercise 4M: Conducting the Session

Now that you have the background and can understand some of the perspectives held by the community members in relation to the land titles issue, it's time to try to conduct the GMB session. Once everyone is gathered together in a room, the facilitator can begin by explaining what the focus of the session will be (in this case, the land titles issue).

In some GMB sessions, the participants may not be familiar with CLDs, and so the notation and meaning of causal relationships would have to be explained first.

Once everyone is clear on the focus of the session, and on how CLDs work, you can begin the modelling. There are many ways to being a GMB session. A list of 'scripts' for how to carry out GMB sessions can be found in this free online resource.

For this session, we recommend that you begin the modelling by drawing up the central variable on the board, which would be 'Access to land titles'. As in the previous exercises in this module, you can work from here to identify what affects and is affected by this variable. The only different is that you are now doing this exercise as a group, rather than an individual. This can be considerably more challenging, but it may also lead to more fruitful insights, as different stakeholders share different knowledge or perspectives about the problem.

As the group begin identifying causal relationships, the modeller should begin drawing these on the board/computer. When necessary, the modeller should clarify whether she/he is modelling the thoughts of the participants correctly. The modeller should also look out for feedback loops in the model. When one is identified, the facilitator can ask the group to discuss whether this is a reinforcing or balancing loop. If it is reinforcing, the facilitator should ask whether it is currently working in a desirable or undesirable direction. This exercise can help participants move from linear to non-linear reasoning and insights about the system at hand.

Exercise 4N: Identifying Solutions

Once you have constructed a model of the land titles issue, the facilitator could begin to ask the group to use the model to identify leverage points, i.e. variables on which they have control and which can lead to a desirable change in the system (in this case, the actions that could help in getting the land titles). It's important that the facilitator remains neutral in this and at all stages of the session. The facilitator's role is to manage group dynamics and to keep the discussion within the range of focus, but they should not let their own opinions influence the group's decisions or beliefs.

Typically, once a group has identified potential solutions to the issue at hand, the facilitator can conclude by reflecting on what has been done in the session, and what the decided outcomes were (in the case of Martina Bustos, the group focused on identifying how to resolve the land titles issue). Sometimes the outcome of a GMB session can be purely educational, but ideally a GMB session will lead the group to identify concrete steps or strategies that they can enact in order to resolve a common problem.

There is no right or wrong way to model the issues that the Martina Bustos community spoke of. Nonetheless, if you are interested in seeing the way we chose to model it, then you can see the model in section three of the next module, Module Five, which focuses on case studies in which CLDs were used.

Trying Another GMB Session

Having completed this session, you may wish to build on your experience and conduct another GMB session with your group, this time on an issue of your choosing. Ideally, the topic you choose would be something that everyone is interested and invested in, and something that would be interesting to study using system dynamics. This would ideally be something that involves a complex system with many interacting parts and the potential for unintended consequences.

End of Section 4: A First Experience of Group Model Building

Facilitating a GMB session is not an easy task, nor is gathering a group of people to participate in one. Nonetheless, both the facilitation and the gathering of people are real life challenges that you may face should you decide to practice system dynamics, and so we wanted to give you an opportunity to challenge yourself in this way. If things did not go smoothly, or as planned, that is more than understandable and part of the learning process. The goal is to achieve incremental improvements every time you facilitate, or even participate, in a GMB session.

Section 5: Tips for Model Building

The previous sections have led us through several cases where you have built your own models. We hope that by now you are beginning to feel comfortable in constructing your own causal loop diagrams. In this section we will step back and reflect on some of the best practices for model building.

Modelling Tips

As we mentioned at the beginning of this module, there are many different processes by which a model can be built, and our recommended process is just one. As you gain experience building models, you are likely to begin doing so naturally, with less focus on the process and more focus on the issue.

Nonetheless, you may find the below list a useful resource as a rough step by step guide for one way of developing a CLD:

- Collect your data (i.e. do interviews with relevant stakeholders/experts, read articles, collective quantitative data, etc.)*.
- Determine the boundary of your model: what issues you will include, and what ones you will leave out. Note that this is an iterative process.
- Determine the resolution of your model: decide on the right level of abstraction/detail. This is also an iterative process.
- Start your diagram from your central variable (i.e. the most problematic or important variable).
- Find out what affects and is affected by your central variable, keeping in mind the boundaries that you have set.
- Draw the correct causal links to and from the central variable, denoting whether it is a positive or a negative relationship.
- Highlight significant delays between causes and effects by using a dashed arrow for such relationships.
- Look for and identify important feedback loops in your model, and denote them as either balancing or reinforcing. You may also want to give certain feedback loops a name.

After development we recommend that you:

- Challenge and critically assess the causal links and model boundary/resolution and generally seek to improve your model by iteration.
- Try and determine whether the reinforcing loops in your model are currently acting as a virtuous cycle or as a vicious cycle
- Determine whether the balancing loops are hindering progress or are giving a desired stability to the system
- Identify the major assumptions in your model and make these assumptions clear to others (for example if you write a paper about a CLD, we recommend having a section called "model assumptions" which explains each assumption and ideally gives evidence to support each one)

- Create a dynamic story of the current problem, unfolding your model step by step

Regarding the layout of variables on a page/screen, we recommend that you:

- Try to make loops look like loops! In other words, use curved lines for feedback loops so that they look like circles or ovals on the page, making them more easily identifiable
- Organize your diagrams to minimize crossed lines.

Perhaps the most important part of the model building process is to express any insights or clarity you have gained as a result of the process. This kind of reflection can help you to see the value and the limitations of CLDs.

---

*The information on which you base your CLD can come from many sources. This includes surveys, organisational databases, news papers, research papers, government publications, TV interviews, industry statistics, video recordings, transcripts, personal interviews, personal observations, structured observation, analysis and experimentation, and much more.

When Can We Use CLDs?

Before we conclude this module, it's important that we reflect a bit on the use of CLDs. Of course, CLDs are useful tools in many contexts, but not all. As the saying goes, having a hammer does not mean that everything is a nail!

So when are CLDs likely to be useful?

There is no hard and fast rule for this, as it depends on judgement. However, below are some scenarios in which CLDs have proven useful:

In analysing problems where a system or variable within the system is changing or suspected to eventually change in an undesirable way. For example, if the number of high-school drop outs is increasing, the inequality in a country is on the rise, or if a population of a specific species is declining, CLDs can be useful in developing and communicating an understanding of why.
CLDs can also be used when a group of people hold differing views on what is causing a certain problem, and hold different knowledge about that problem which requires synthesis in order to bring about a more holistic understanding. For example, people from different departments in a business might hold different views about why the business is losing market share, and a CLD can be a very useful tool for synthesising these different perspectives and improving the likelihood of finding a successful solution. It also serves as a great communication tool between different stakeholders in such settings.

Limitations of CLDs

But what then, are the limitations of CLDs? We briefly discussed this in Module Two, but it may be worth revisiting now that you have a deeper understanding of what CLDs are.

Firstly, as systems thinkers we are interested in understanding how things change over time. CLDs, however, hold a limited ability to provide an understanding of the behavioural output of a system. Without quantifying your model, it is difficult to infer the magnitude of change in different variables over time, what feedback loops are dominating the system at any given moment, and when shifts in feedback dominance are likely to occur. In order to better understand the behaviour created by the system structure, a formal simulation model is needed.

Another limitation of CLDs is that it is sometimes difficult to balance a need to simplify with a need to be precise. When simplifying in order to shed light on a specific issue or part of the system, the underlying rationale or the assumptions behind the causal links included in the CLD can become somewhat difficult to grasp. This can be particularly true when one has not been involved in the creation of the CLD. Whilst the meaning of all variables might be clear to those who created the model, the meaning to those reading about it may not be so clear. This limitation can also be true for simulation models.

For further readings on this topic, see this paper by David Lane.

End of Section 5: Tips for Model Building

That brings us to the end of section 5. This offered us a chance to reflect on some of the practices we have learned so far, and on the reasons for which we use CLDs in the first place, as well as some of their limitations.

In the next section we will introduce a new concept from system dynamics: that of stocks and flows.

Section 6: Identifying Stocks and Flows

Up until this point in the course, we have built CLDs where there is no distinction between different types of variables. However, system dynamicists would argue that it is often important to make a clear distinction between two types of variables: stock variables (simply known as stocks), and flow variables (known as flows). In this section we will discuss the difference between stocks and flows, and why it is so important to differentiate between them.

We will also demonstrate the notation used to highlight and differentiate stocks and flows in a system dynamics model. Differentiating between stocks and flows is typically done in simulation models rather than in CLDs, but some CLDs also make the distinction. The importance of making this distinction should be clear by the end of this section.

Exercise 4O: The Department Store Task

To begin this section, it will be worthwhile to do a small exercise, developed by professor John Sterman of MIT. The exercise asks you to review the below graph, which shows the number of people entering and leaving a department store over a 30-minute period, and then answer the questions found below. If you believe the answer can't be determined, then simply leave the answer blank.

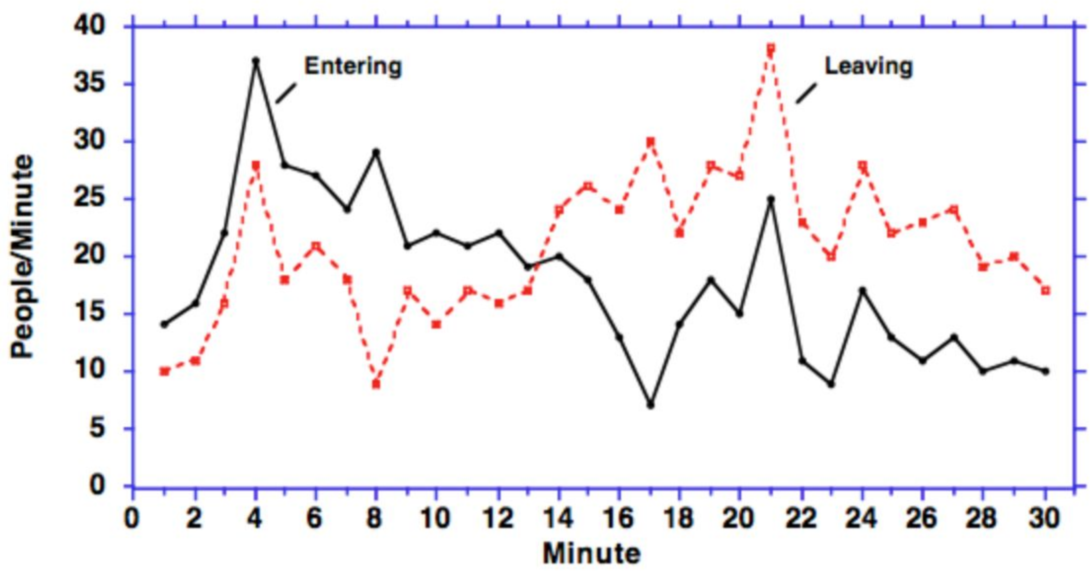

Questions

1. During which minute did most people enter the store?
2. During which minute did most people leave the store?
3. During which minute were the most people in the store?

4. During which minute were the fewest people in the store?

Answers to 4O

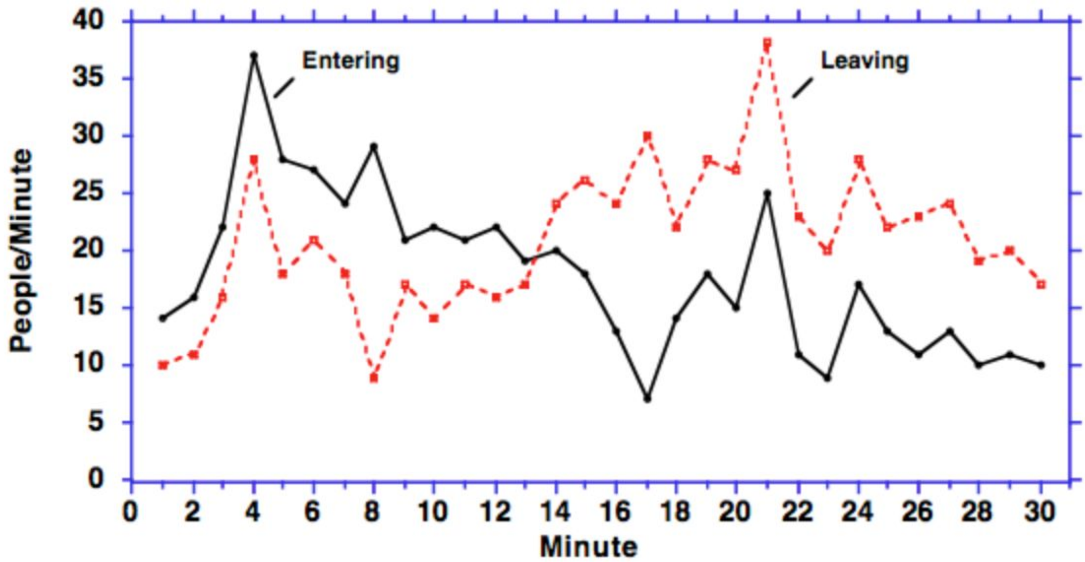

The answers are as follows:

1. Most people entered the store at minute 4.
2. Most people left the store at minute 21.
3. The most people were in the store during minute 13.
4. The least people were in the store during minute 30.

Reflecting on the Exercise

You may have answered Questions 1 and 2 correctly. In this study conducted on MIT students, 96% of participants answered Question 1 (about when most people entered the store) correctly, and 95% of participants also answered Question 2 (about when most people left the store) correctly.

However, you may have had more trouble with questions 3 and 4 (about when the most and least people are in the store, respectively). In the above-mentioned study, only 44% of respondents answered Question 3 correctly, and just 31% of respondents answered Question 4 correctly.

For now, let's focus on the responses to Question 3. 29% of respondents to this question incorrectly asserted that most people are in the store at minute 8, which is the minute during which there is the greatest difference between the people entering and leaving the store. Many people incorrectly assume that this means that this will be the minute during which the most people will be in the store. In reality, the number of people in the store will continue to grow for as long as there are more people entering the store than leaving the store. For that reason, we know that most people are in the store at minute 21, as up until then the number of people entering the store is greater than the number of people leaving. After minute 21, there are more people leaving than entering, which means that the number of people in the store starts to decrease.

Accumulation

The reason many people fail on questions 3 is because they fail to understand the principles of accumulation. In system dynamics terms, we would say that they fail to distinguish between stocks and flows. Flows are variables that happen over time: for example, the amount of water flowing from a tap per second is a flow, and the water flowing down the drain per second is also a flow. Stocks refer to variables that accumulate flows. For example, the level of water in a bathtub is a stock that accumulates the flow of water into and out of the bathtub. The water flowing into the bathtub per second would be known as an inflow to that stock (because it adds to the stock) whereas the water flowing out of the bathtub per second would be known as an outflow from that stock (because it subtracts from it).

Similarly, the number of people in a department store is a stock, and it accumulates the flows of people entering the store per minute (the inflow), and people leaving the store per minute (the outflow). The important thing to remember about stocks and flows is that a stock will increase if and only if its inflow(s) is greater than its outflow(s). So, the number of people in a department store will always increase when the people entering the store is greater than the people leaving the store.

Notation for Stocks and Flows

Typically, CLDs do not distinguish between stock and flow variables. For example, in the department store task, a simple CLD of the situation might look like this:

However, if we wanted to highlight the stocks and flows in this model, then we would use notation as seen below:

The rectangular box around the stock represents the fact that this is a variable where things accumulate or drain away from. The hourglass symbol for the flows symbolise the fact that they always happen over time (per second, per minute, per month, per year). Note that the inflow (people entering the store per minute) is pointing into the stock, whereas the outflow (people leaving the store per minute) is coming out of the stock. This represents the fact that inflows add to a stock, whereas outflows subtract from a stock.

Determine Stocks from Flows

So the easiest way to determine if a variable is a stock or a flow is to ask yourself whether the variable happens over time (i.e. per second, per week, per year). If it does, then it is a flow. And if the variable is something that accumulates flows, then it is a stock.

Exercise 4P: Recognising Stocks and Flows

The distinction between stocks and flows might not be immediately obvious to you. However it should become clearer when we begin to consider more and more examples of what a stock is, and what a flow is.

Below you can see a list of variables. Can you spot the stocks and the flows? Some of these stocks and flows are causally related. Can you also draw the stock and flow relationships between these variables, using the notation presented on the previous page? Remember that inflows are arrows pointing into the stock, and outflows are arrows pointing out from the stock.

The answer is on the next page.

Variables

- Water in a bathtub
- Litres of water flowing into a bathtub per second
- Litres of water flowing out of a bathtub per second
- World population
- Deaths per year
- Births per year
- Money in a bank account
- Dollars withdrawn from a bank account per month
- Dollars deposited into a bank account per month
- Goods in a warehouse
- Products delivered to a warehouse per day
- Products shipped from a warehouse per day

- Houses in a neighbourhood
- Houses constructed per month
- Houses demolished per month
- Trees in a forest
- Trees dying/cut per month
- Trees grown/planted per month

Answers to 4P

Did you recognise the stocks from the flows? And did you draw the draw the inflows and outflows correctly?

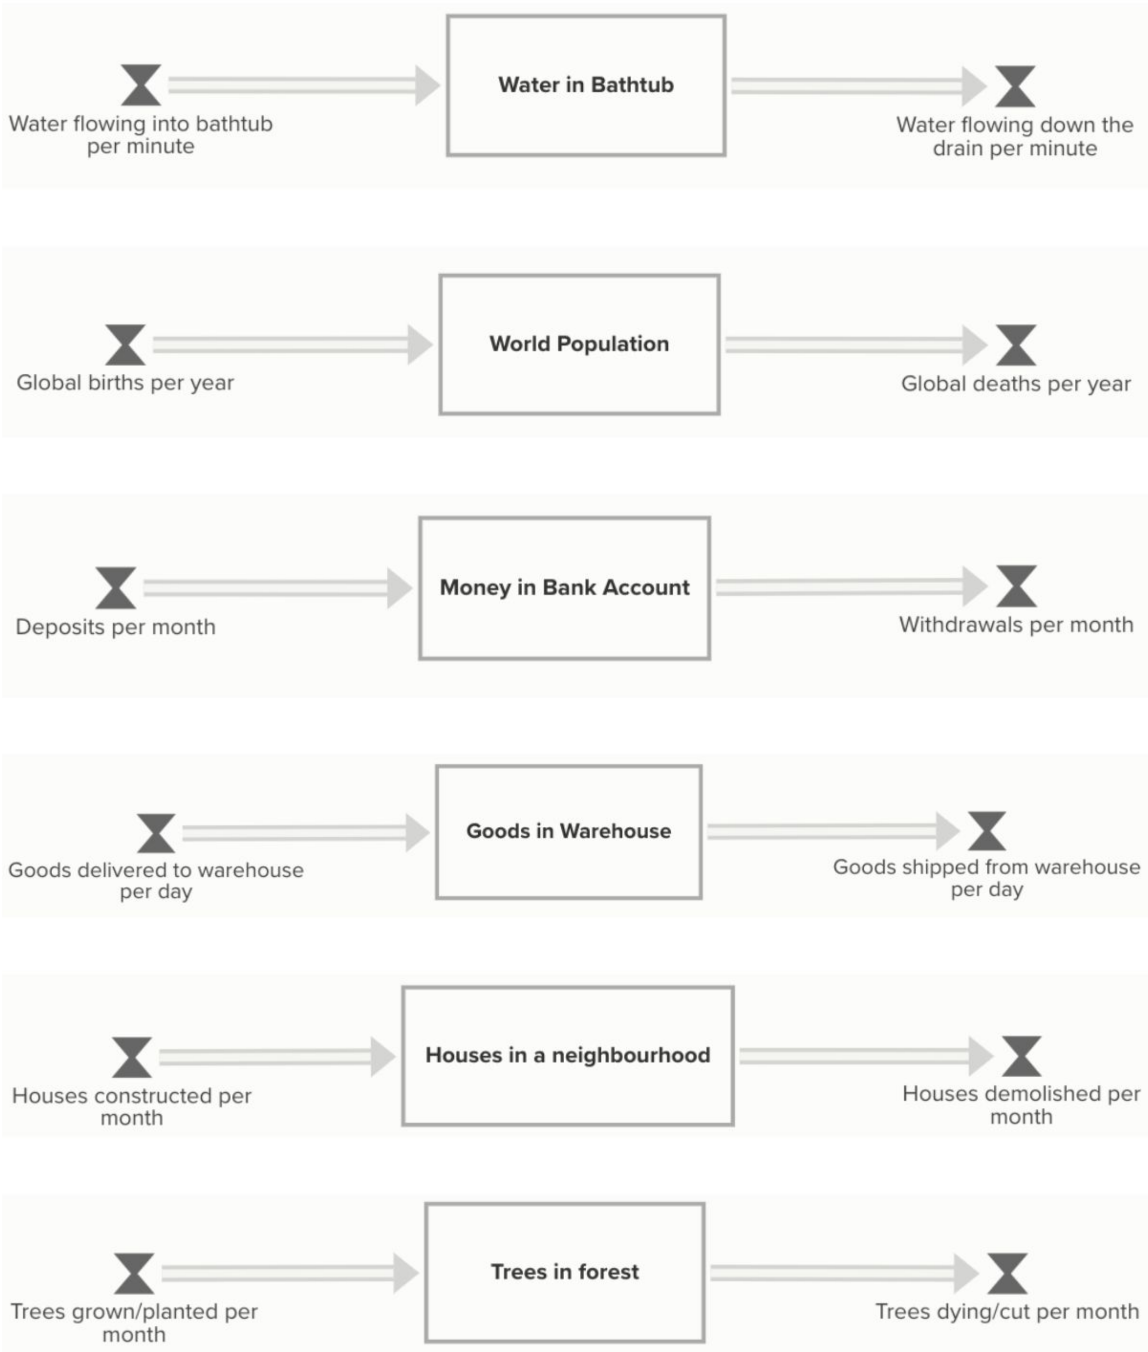

Not All Stocks Are Material

Note that stocks do not always have to refer to physical things. Intangible variables can also be stocks. For example, if you were building a model about the reputation of a certain brand, you might highlight the fact that reputation is something that accumulates over time. As such you might model it as a stock, which is affected by flows such as "positive experiences with the brand per month" (an inflow), and "negative experiences with the brand" (an outflow).

The Importance of Distinguishing Stocks from Flows

The research on the department store task demonstrates that many people have trouble with the principles of accumulation, i.e. with distinguishing stocks from flows. In the context of determining the number of people in a store at any given time, the common inability to recognise stocks form flows is not a life or death situation. Nonetheless, this cognitive blind spot can have important consequences. For example, many studies have shown that people don't understand this fundamental characteristic of the greenhouse gas (GHG) problem (for example, see this study). Indeed, it is often assumed that if we stabilise our GHG emissions, then the stock of GHGs in the atmosphere will also stabilise. But is this true? Given that our current rate of GHG emissions is far above the rate at which GHGs are absorbed by natural sinks, what do you think will happen if we stabilise emissions at their current rate?

Write down your answer now, and we'll come back to it after doing the next exercise.

Exercise 4Q

To help you think through this, try to draw a stock and flow model of how the level of CO2 (the main greenhouse gas) in the atmosphere changes. To keep things simple, you should consider just three variables:

- The CO2 emitted into the atmosphere per year
- The amount of CO2 in the atmosphere
- The CO2 absorbed from the atmosphere each year (through carbon capture from trees and the seas, for example).

The answer is on the next page.

Answer to 4Q (1/3)

The amount of CO2 in the atmosphere is a stock, like the water in the bathtub.

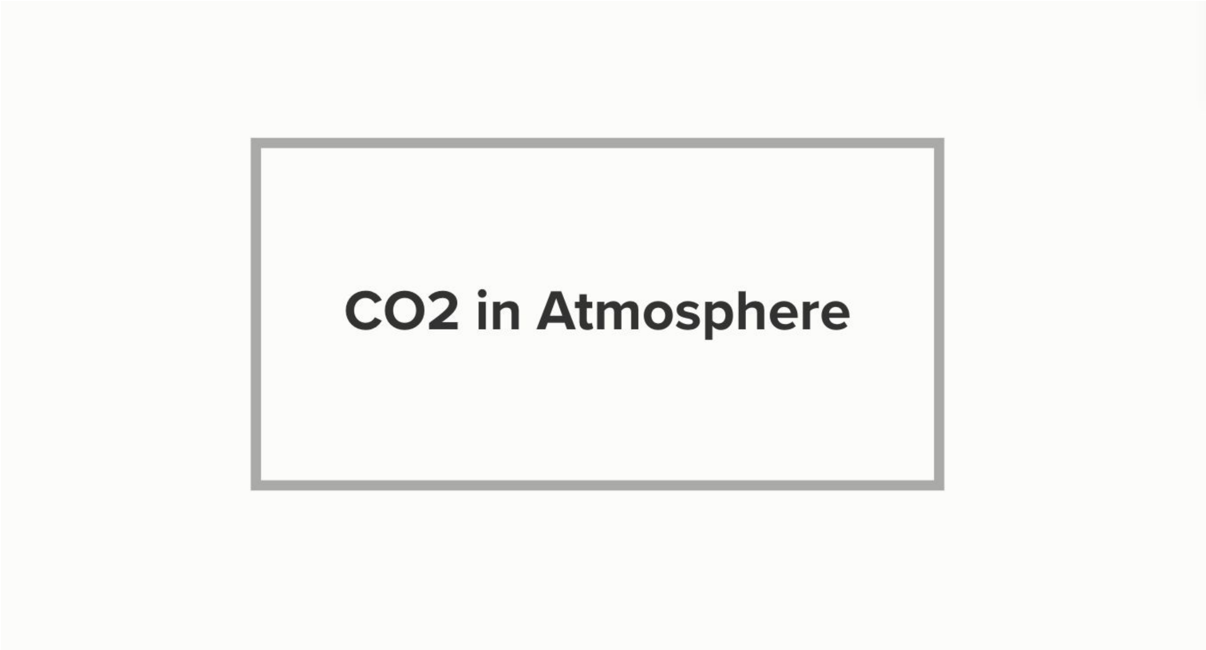

Answer to 4Q (2/3)

The annual global CO2 emissions is the inflow to this stock, like the water flowing through the tap into the bathtub.

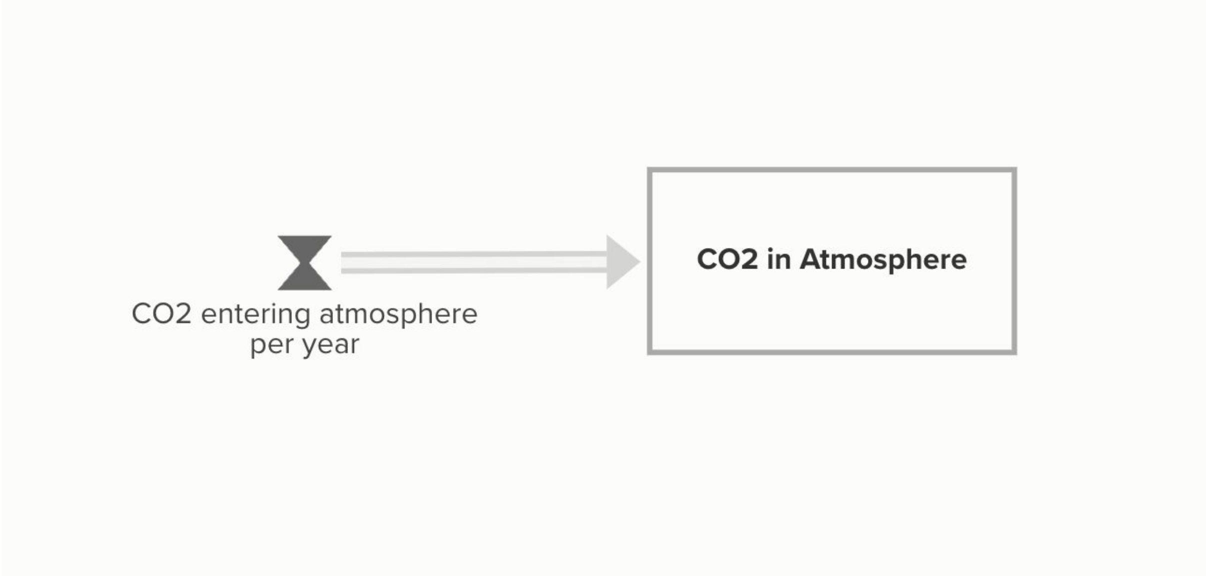

Answer to 4Q (3/3)

Finally, the annual CO_2 absorbed from the atmosphere is the outflow, like the water flowing down the drain.

This removal of CO_2 from the atmosphere happens through natural processes. For example, when trees photosynthesise, they take CO_2 out of the atmosphere and use it to grow. Ecosystems that take carbon out of the atmosphere are known as carbon sinks.

Did you correctly distinguish the stock from the flows, and the inflow from the outflow?

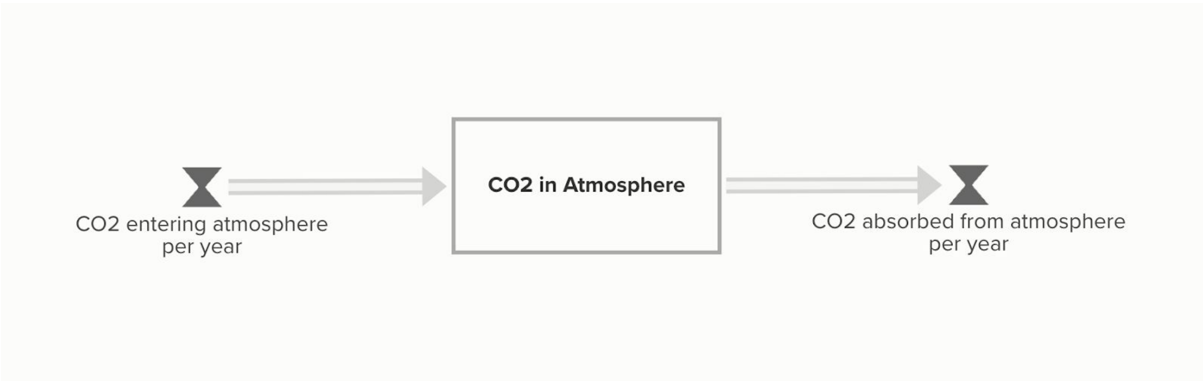

Exercise 4R

Given that we now understand the GHG problem in terms of stocks and flows, what do you think would happen if we stabilised CO_2 emissions at their current rate, given that they are currently far higher than the rate of CO_2 being absorbed from the atmosphere? Would the stock of CO_2 in the atmosphere continue to grow, stabilise, or go into decline? Has building the stock and flow model of the situation has helped you think through this issue in any way?

Answer to 4R

If we stabilise emissions at current rates, then the CO_2 in the atmosphere will not stabilise, but will continue to grow. In order to stabilise the level of emissions in the atmosphere, it is necessary to reduce the rate of emissions to the rate of CO_2 being absorbed. In other words, we would need to reduce the inflow to the same level as the outflow. In order to reduce the level of CO_2 in the atmosphere, we would need to bring the emissions to lower than the rate of absorption.

This can be summarised in the below graph, which displays CO_2 concentration levels in the atmosphere under different emission scenarios.

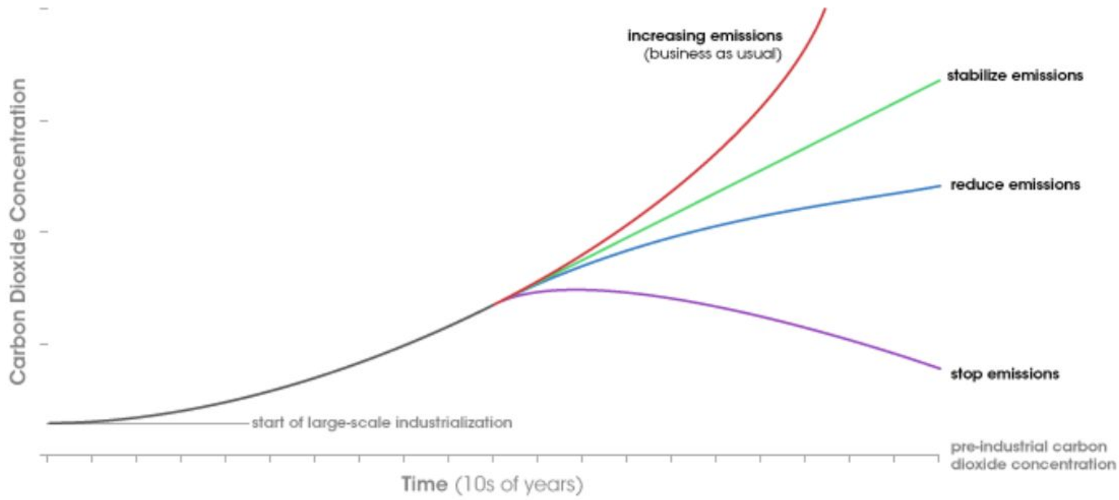

Image courtesy of NASA Earth Observatory.

Visualising the GHG Problem

To understand how flows affect stocks, it can often help to visualise it. The simulator available on this webpage* is based on a quantified system dynamics model and allows you to see how different CO2 emission pathways will affect the stock of CO2 in the atmosphere. Besides increasing our understanding of the greenhouse gas problem, this tool also offers a great way to visualise how stock and flow systems work. Try and run the simulation under different emission scenarios and see how that affects the level of CO2 in the atmosphere over time.

**If the simulator doesn't work at first, try using another browser and/or check your Adobe Flash Player settings.*

End of Section 6: Identifying Stocks and Flows

That brings us to the end section 6. Being able to differentiate between stocks and flows is essential when building simulation models, but it is not necessary when creating CLDs. Nonetheless, as we hoped to have shown, the distinction is still an important one in general and can help improve reasoning about complex issues – such as climate change and the accumulation of GHG emissions in the atmosphere

End of Module Four

That brings this module to a close. To recap, we began with an exercise building a model of a fairly simple and relatable topic, based on an imagined interview with two university students. From here, we moved to modelling a more complex topic (the utility death spiral), where we hope to have shown how causal loop diagramming can help in both thinking more clearly and analytically about an issue, and in identifying the assumptions implicit in someone's mental model. The third exercise, on the Theory of the Invisible Hand, provided a greater challenge in asking you to build a CLD based on a video, and also demonstrated how CLDs can be used to analyse and hopefully better understand abstract theories. In section four we asked you to imagine yourselves in a real-life group model building situation, a challenging task no doubt. In section five, we reflected on the best practices for creating CLDs, and touched again on its limitations. In section 6, we introduced the concepts of stocks and flows and discussed why it is so important to differentiate between them.

Further readings

Sterman, J. D., & Sterman, J. D. (2014). Business dynamics: Systems thinking and modeling for a complex world. Boston: Irwin. (we recommend chapters 1, 3, and 5)

Homer, J. B. (1996), Why we iterate: scientific modeling in theory and practice. Syst. Dyn. Rev., 12: 1-19. doi:10.1002/(SICI)1099-1727(199621)12:1<1::AID-SDR93>3.0.CO;2-P

Homer, J. and Oliva, R. (2001), Maps and models in system dynamics: a response to Coyle. Syst. Dyn. Rev., 17: 347-355. doi:10.1002/sdr.224

Kim, D. H., & Kim, D. H. (1995). Guidelines for drawing causal loop diagrams. Waltham, MA: Pegasus Communications.

Sterman, J. (n.d.). Fone Tuning Your Casual Loop Diagram - Part 1. Retrieved from: https://thesystemsthinker.com/fine-tuning-your-causal-loop-diagrams-part-i/

Rushing, W. (n.d.). Causal Loop Diagrams: Little Known Analytical Tool. Retrieved from https://www.isixsigma.com/tools-templates/cause-effect/causal-loop-diagrams-little-known-analytical-tool/

Module Five
Case Studies

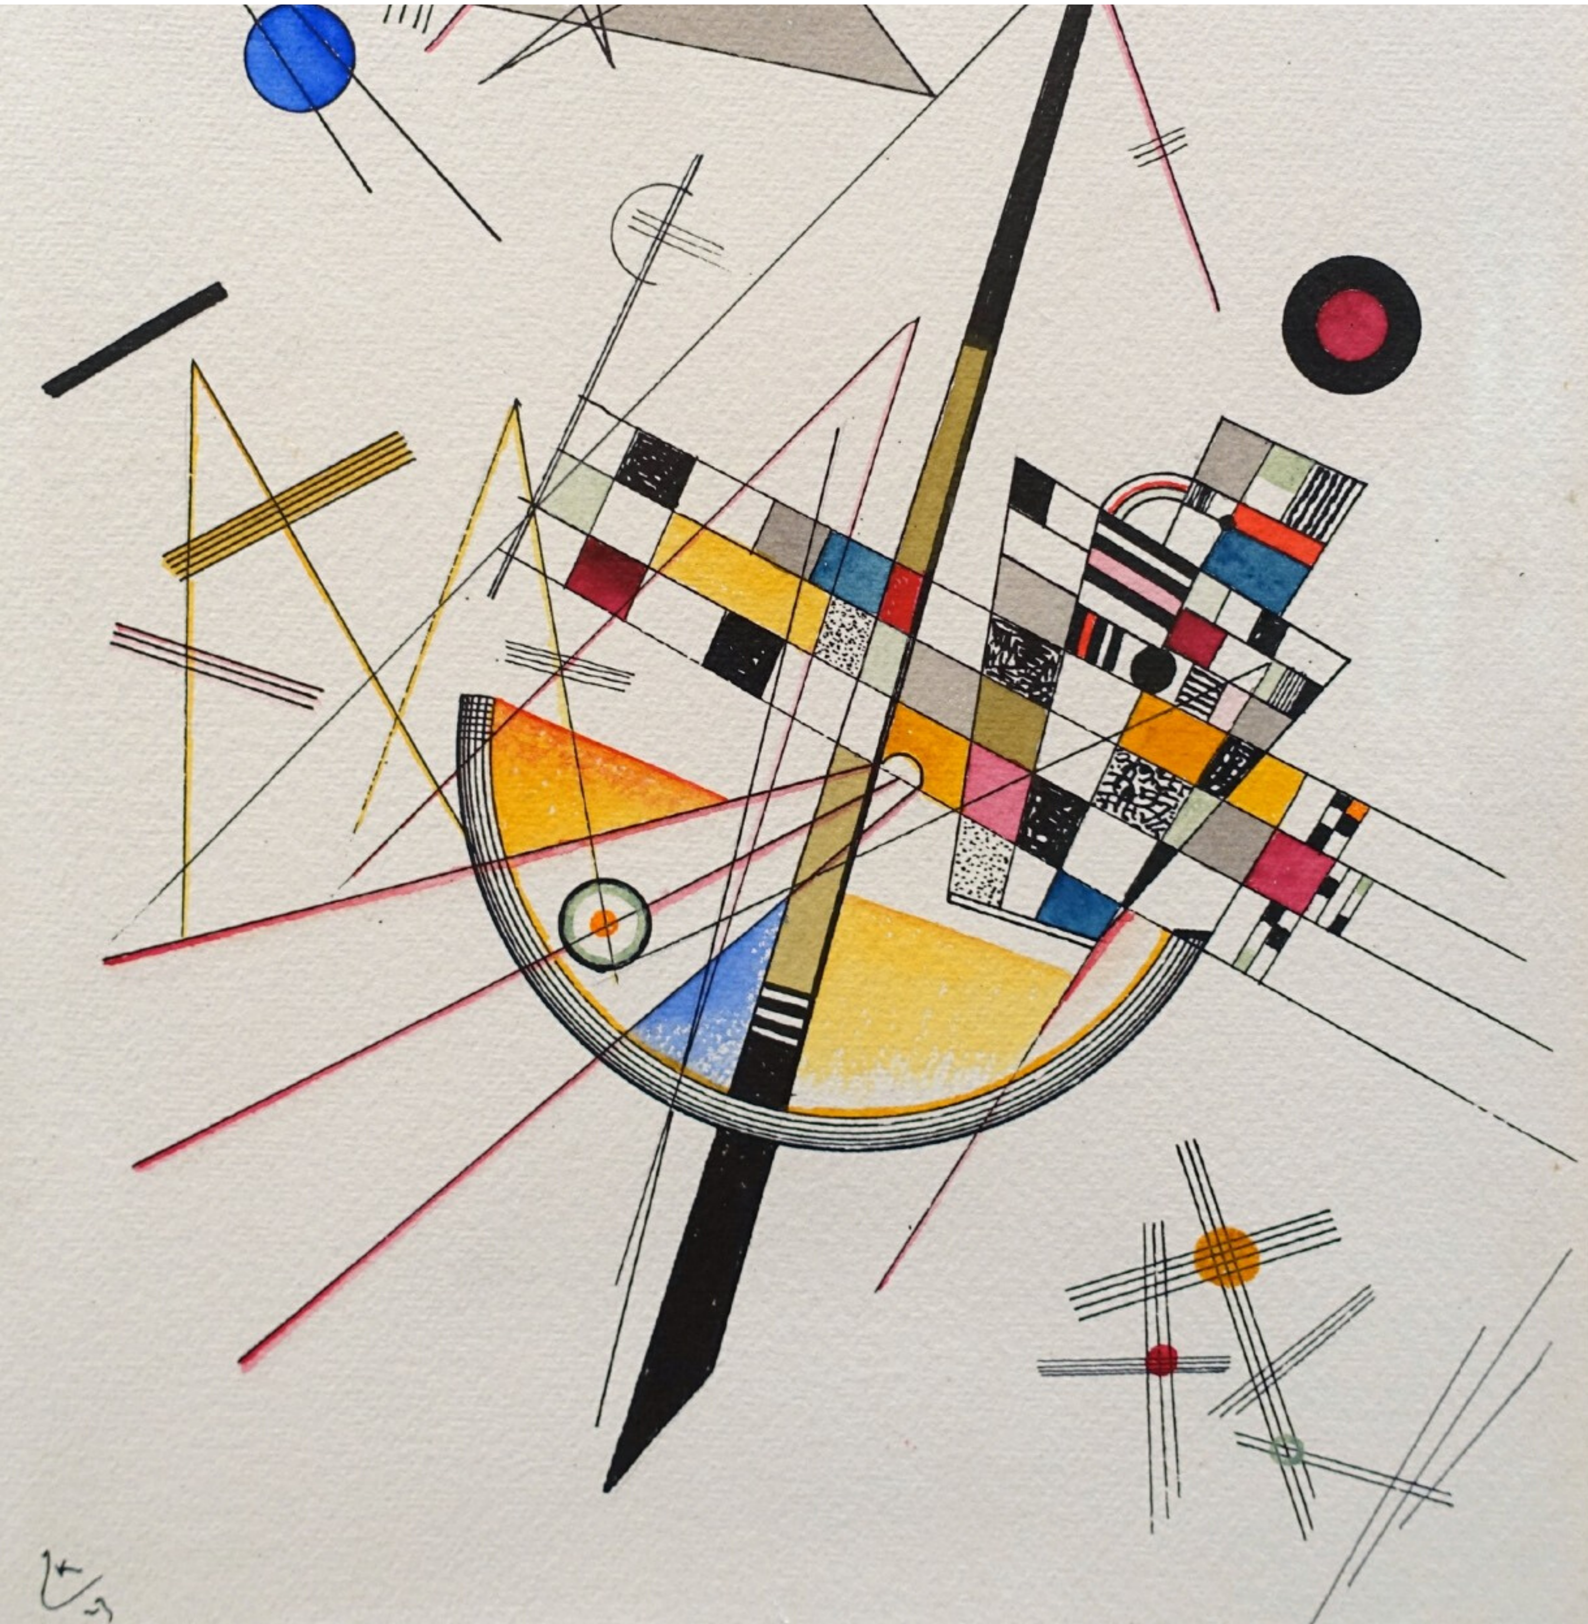

WASSILY KANDINSKI [CC0 OR PUBLIC DOMAIN], FROM WIKIMEDIA COMMONS

Module Five Overview

Introduction

Module Five discusses five case studies that demonstrate how CLDs have been used in certain situations, and the value that they provided. This is then followed by some final thoughts and a final exercise to conclude the course.

The previous module walked you through the process of building models relating to different problems. In this module the focus will be on real life case studies in which causal loop diagrams were used to help analyse a certain problem/highlight a certain issue.

The real value of the process of causal loop diagramming cannot be fully appreciated by analyzing case studies alone. Oftentimes, one needs to be experiencing the complexity of the situation firsthand in order to see how modelling can help you or the group you are a part of to communicate and think clearly, to ask the right questions, and to feel a sense of comfort that you are able to manage the complexity of the situation without shying away from it. Nonetheless, the following case studies should provide a sense of how CLDs can help in certain situations, and will also give you some idea of the wide variety of contexts in which they have been used.

Module Outline

We will deal with five case studies which are quite different in nature. We hope that by the end of this chapter you will have a better understanding of how causal loop diagrams have been used historically, and how you yourself might use them to help tackle a challenge in your own life/organisation/society.

The five case studies are as described below:

1. Gender Equality in Dutch Technical Universities (Inge Bleijenbergh and Marloes Van Engen)

2. Hewlett-Packard Market Share (Deborah Campbell)

3. The Martina Bustos Community (The Sustainability Laboratory)

4. The ISLAND project (Andrea Bassi, Gilbert Probst, & Prakash Deenapanray)

5. Building a Healthcare System (Markus Schwaninger)

Case study 1: Gender Equality in Dutch Technical Universities

The Context

Gender inequality is a complex problem. Indeed, stakeholders often disagree on the existence of the problem itself, its causes, or its solutions. The problem is particularly evident in Dutch technical universities, which have a very low ratio of female to male professors. In addition, the progress towards an equal ratio is considered slow.

The case study we discuss here was performed in the Social Sciences and Humanities University and the Technical University in the Netherlands. This case occurred in 2009, when professors Inge Bleijenbergh and Marloes Van Engen along with their research team used participatory modelling as an intervention methodology to "support stakeholders in: reaching a shared problem definition and analysis of gender inequality; and identifying and implementing policies to tackle gender inequality."

The Process: Participatory Modelling and CLD Development

Participatory modelling is a method in which different stakeholders to a problem get together to create a shared model of the problem. This is typically done in day or half-day long workshops, which are usually facilitated by one or two experts in systems modelling. The aim is to increase understanding of the problem, synthesise information from different perspectives, reach consensus between stakeholders with opposing viewpoints, and ultimately come up with a strategy for solving the problem.

The researchers in this case study used CLDs as a modelling methodology and involved many different stakeholders in various stages of model development. They held in-depth interviews with 44 academics from universities that suffer from low female-to-male ratios at the assistant and associate professor level. The results of these interviews were then discussed with "five focus groups with 34 stakeholders."

Researchers along with stakeholders developed models to represent the problematic issue at hand for each school independently. A generic structure was generated to represent the common structure among these schools that gives rise to the problematic behaviour, i.e., the slow progress towards a more equal ratio between female and male assistant and associate professors. A set of actionable policies were then put forward by the researchers based on insights they gained from the model, which can be seen on the next page.

The Model

A slightly simplified version of the model produced by the team of researchers is shown to the right*. It identifies three main reinforcing feedback loops that go some way to explaining why the ratio of female to male professors is so low in Dutch technical universities.
Note that this CLD identifies stocks and flows, which we discussed in the previous module. In this case the participants felt it was useful to use stocks and flows to identify what is known as an "aging chain" by which women go from being PHDs to associate professors to full professors. Note that many outflows from one stock are an inflow to another.
We can now discuss each loop separately, which will also help us to better understand this stock and flow structure.

*Please note: this model is our own recreation of the research team's model. It has been altered slightly for educational purposes.

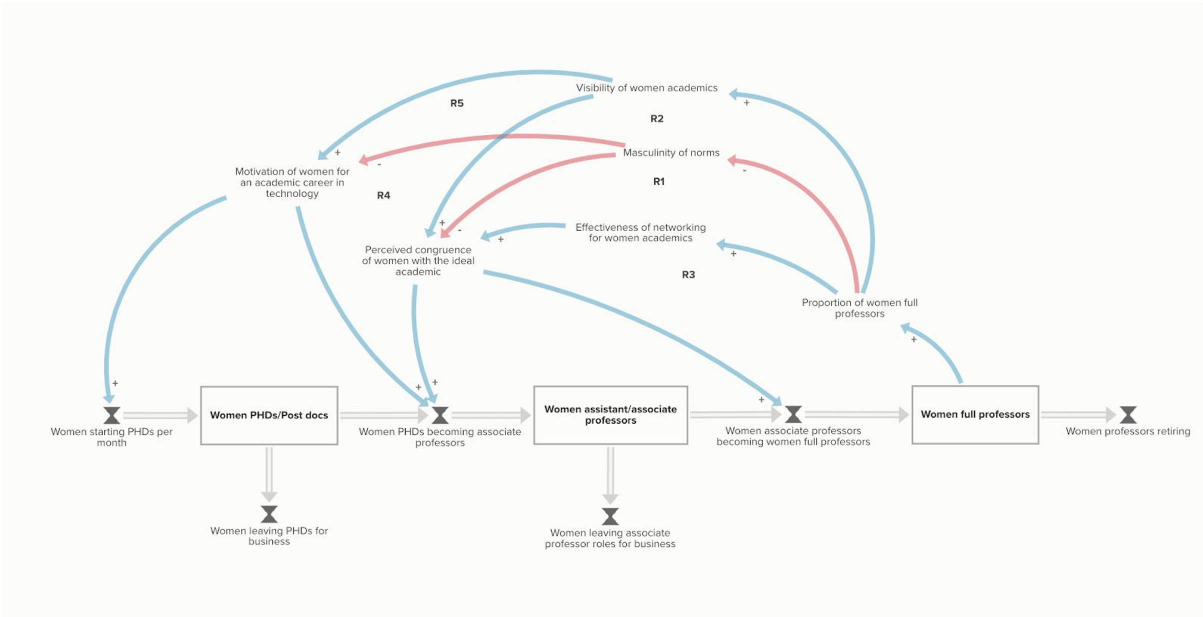

R1: Masculinity of Norms

The reinforcing loop R1 says that the low Proportion of women full professors leads to a Masculinity of norms in Dutch technical universities. This reduces the Perceived congruence of women with the ideal academic which reduces the number of women PHDs deciding to become associate professors. This in turn reduces the number of Women full professors, which further reinforces the Masculinity of norms.

By masculinity of norms, the researchers were referring to the culture of masculinity that existed in each university. Masculine culture had different meanings to different participants, but the researchers were happy enough to have just one variable to represent this broad concept in this high level model.

The basis for the importance of recognising the potential masculinity of norms came from the researchers' interviews. For example, consider the following excerpt from an interview with a male Dean at the Technology institute.

The researcher asked:

"To what extent do you think your policy creates unintended obstacles to hiring, selecting and promoting women?"

The Dean responded:

"Culture is the only thing that springs to mind. The culture of material engineering, even more so the culture of maritime engineering, it's quite no nonsense. It does have its positive sides […] but […] um […] it's still a masculine culture."

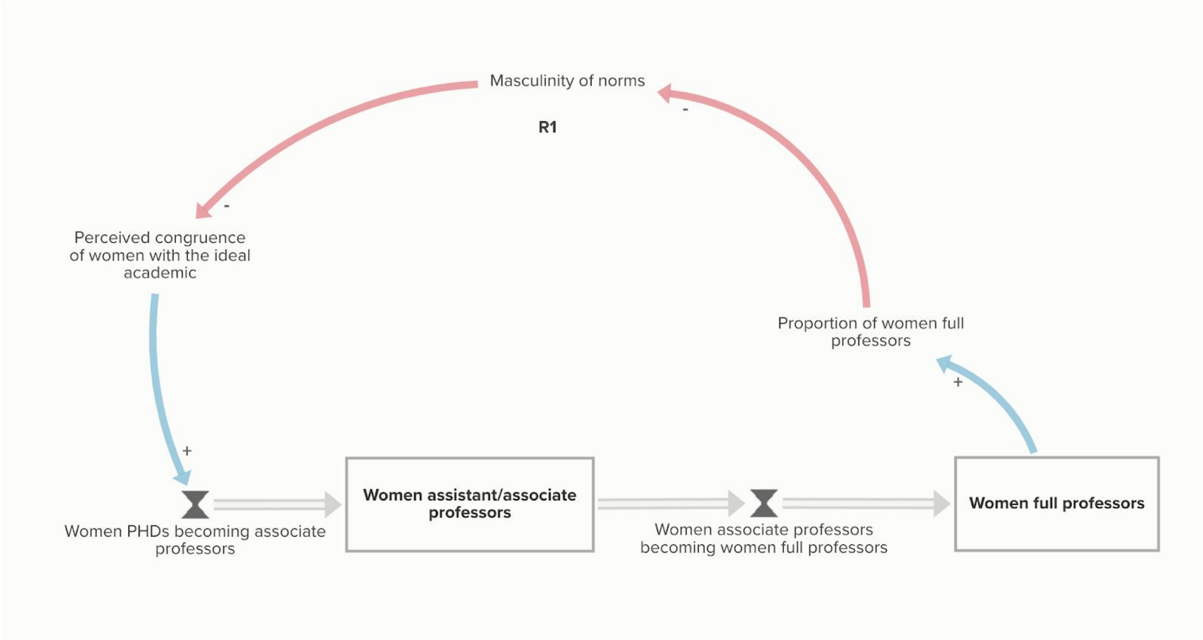

R2: Visibility of Women

Reinforcing loop R2 was also identified as working as a vicious cycle. The low Visibility of women academics meant that there is less Perceived congruence of women with ideal academic role. This reduces the number of women going for associate and full professor roles, which further reduces the Visibility of women academics.

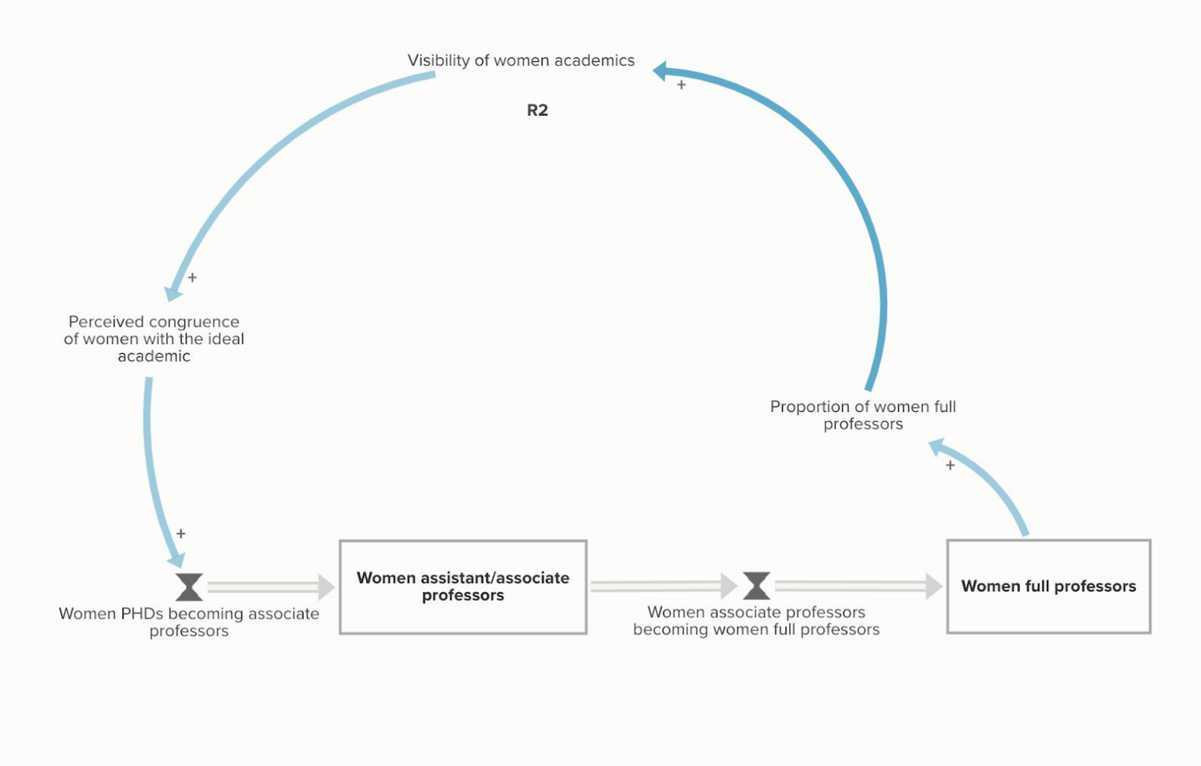

R3: Effectiveness of Networking for Women Academics

The researchers highlighted the importance of networking in their model. During their interviews, the subject of the importance of networking was brought up several times. For example, consider the below two excerpts:

"Sure, you can see all kinds of alliances appearing and, if women are on the Board, they'll make sure that women will get into other positions." (Assistant Professor at the Social Sciences and
Humanities University, woman).

"Traditionally it's a bit of a man's world, so networking and that sort of thing is traditionally easier [for men]" (Assistant Professor at Technology University, man).

The researchers used data such as this to inform the creation of another feedback loop that highlighted how having a higher Proportion of women full professors increased the Effectiveness of networking for women academics, which in turn helped increase the number of Women full professors and so further increased the Effectiveness of networking for women academics.

As with all reinforcing loops, this loop can work as either a virtuous or vicious cycle. This is seen in the perspectives expressed by the interviewees quoted above. Whilst the first quote discusses the potential for women to bring more women into senior academic roles, the second quote points out how the existence of a "man's world" can make networking harder for women at the moment.

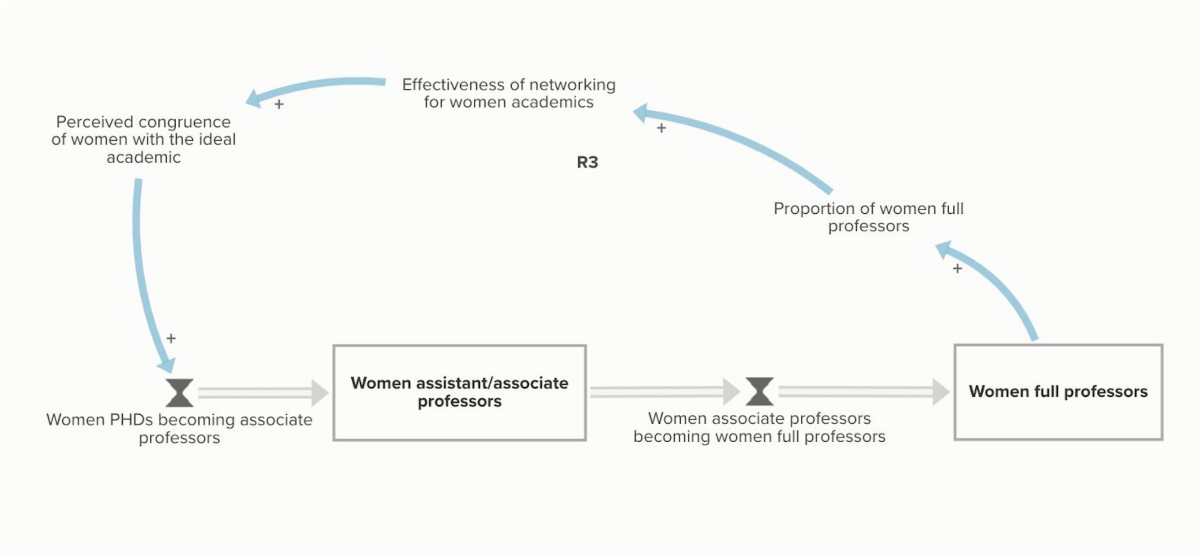

R4: Masculine Culture and Motivation

Another loop identified by the researchers concerns the motivation for women to have an academic career in technology. They identified this lack of motivation as a reinforcing loop, as shown to the right. This loop shows how the Masculinity of norms in Dutch technical universities reduces the Motivation of women for an academic career in technology. This reduces the number of Women full professors, which in turn increases the Masculinity of norms further, thereby further demotivating women to go for more senior roles.

The above describes the loop acting as a vicious cycle. However it can also act as a virtuous cycle, when more women get promoted and so reduce the Masculinity of norms, and thereby increase the Motivation of women to have an academic career in technology.

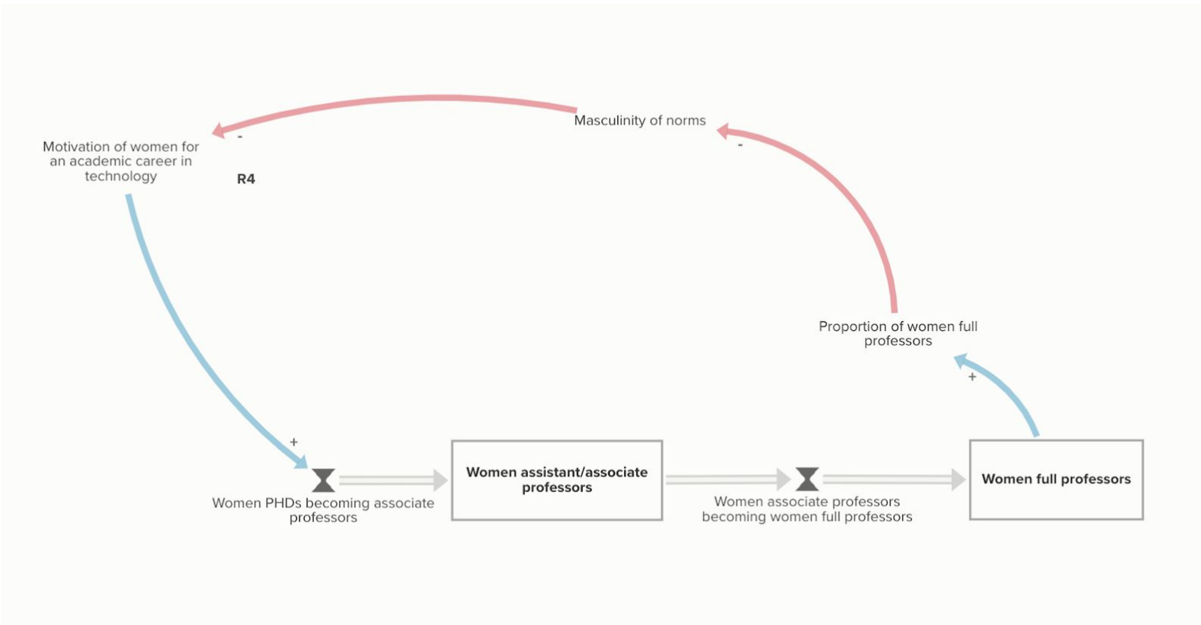

R5: Visibility and Motivation

Similarly, the Visibility of women academics is another way in which women can become increasingly motivated or demotivated to have an academic career in technology. The more visible women are in academia, the more motivated those doing PHDs will be to further their career in this area. Similarly, the less visible they are, the less motivated other women will be.

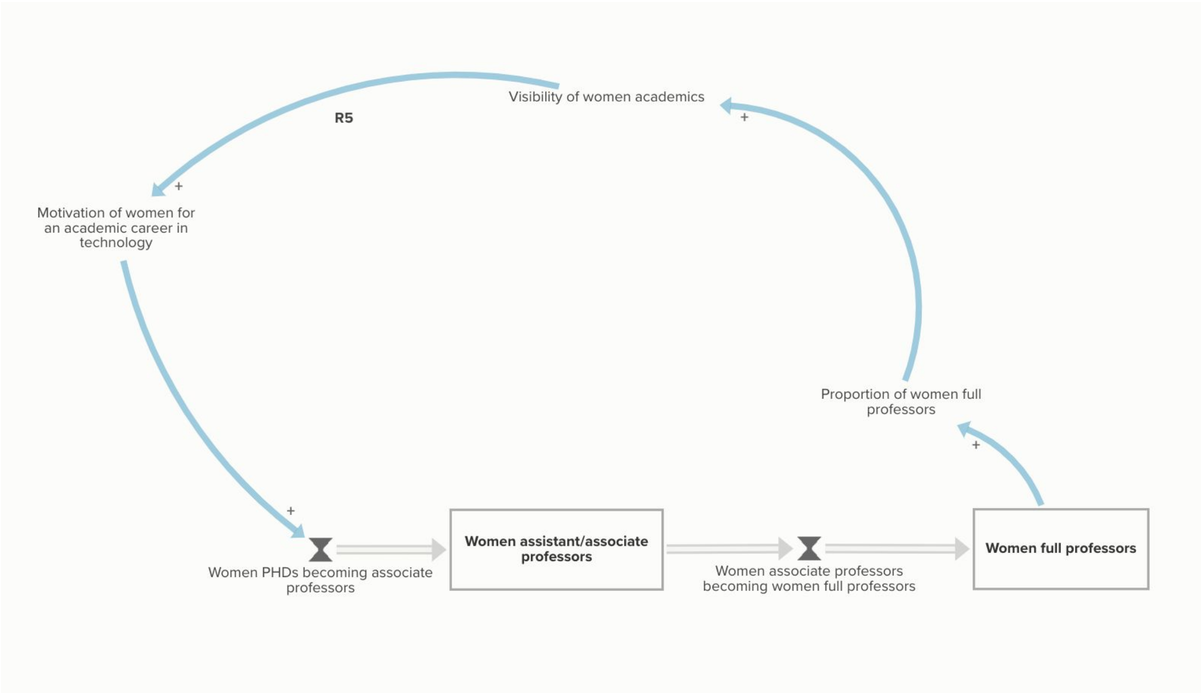

By looking again at the whole model, we can see how all of these reinforcing loops interact. The researchers found that the CLD helped in visualising the interaction between these different causal structures, and highlighted how many feedback loops working as a vicious spiral could be made to work as virtuous cycles.

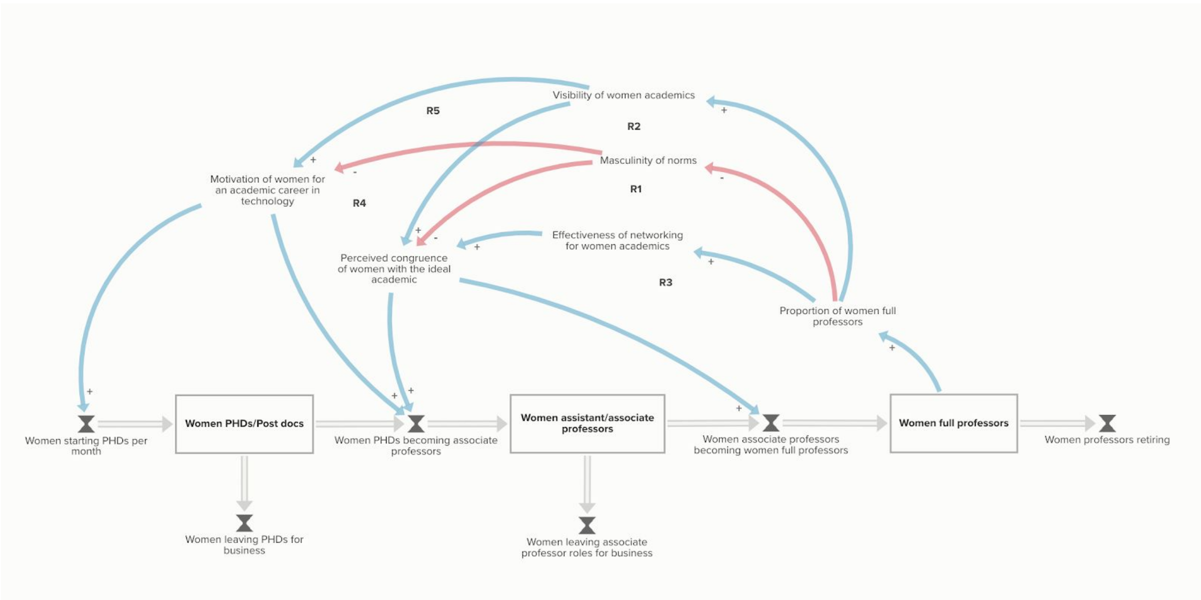

Outcomes: Achieving Consensus Around Actionable Policies

The creation of the CLD in this case was said to have helped in creating a shared understanding of the issue at hand. The fact that it was done in a participatory way was also very important for capturing the perspectives of different stakeholders, and ensuring that they felt that they had been listened to.

Furthermore, the identifications of interrelated feedback processes helped stakeholders identify different intervention policies to solve the issue. The following is a list of some of the policies generated and implemented as a result of the overall process:

"Concerning masculinity of norms, the Social Sciences and Humanities University introduced a correction of publication productivity targets for those taking care leave and part-timers, and for paid research sabbaticals following maternity leave. The Technology University introduced the paid research sabbatical following maternity leave, and relieved the prohibition of flexible positioning in the schemes."
"Concerning visibility of women, the Social Sciences and Humanities University installed a task force that monitors the visibility of women researchers in internal and external communication. Since 2011, selection committees (containing at least one woman) at the Technology University have had to put forward at least three highly qualified women in the discipline on the shortlist of possible candidates for full professors."

Reflecting on Case Study 1

Overall, the researchers in this case believed that they achieved a lot in their research, and that the use of CLDs helped them identify important systemic structures that are inhibiting women's progression to more senior roles in Dutch technical Universities.

The researchers mentioned how "integrating the knowledge of researchers and stakeholders in a causal loop diagram supported the learning of stakeholders about the issue of gender inequality." It also "supported validation of the analysis with organisational stakeholders and, as a consequence, ensured their commitment to the results. Both universities identified and implemented a considerable number of gender equality interventions following the participatory modelling process."

Case Study 2: Hewlett-Packard Market Share

The Context

Hewlett-Packard (HP) is a leading manufacturer and service provider for many consumer products. However, managers at the firm were becoming increasingly concerned about a decline in market share for the company's products. Although HP covered high market share for a specific product, the company was facing stiff competition from competitors offering the same product at lower prices, and incentivising retailers to sell their product.

Deborah Campbell, a senior consultant experienced with systems modelling, suggested participatory modelling* as a possible intervention methodology to understand the overall system and to create a shared understanding of the problem and its possible solutions.

**participatory modelling is the broader term for all kinds of modelling done as a group. Group model building is one school of modelling within participatory modelling.*

The Process: Model Formation and Policy Generation

Deborah and the management team decided to form a team of 15-20 people from the business to develop a CLD as well as a formal in-depth quantitative model. The objective of the exercise was to create a shared understanding of the underlying reasons for the loss in market share and, with that understanding, to develop a strategic action plan to address the problem.

The members of the team came from many different departments within HP, which is important if a systemic understanding is to be gained. Deborah explains that the process was not easy, involving a lot of confusion and uncertainty at different stages of the model's development. Some of the staff, inexperienced with systems modelling, pushed for too much detail in the model, which can often hinder the model's ability to generate high-level insights.

Working Through Confusion

However, Deborah observed that confusion and uncertainty is a natural part of model development. Having the team go through the process of model development, including each department’s role in detail, and then passing through that confusion to reach a new shared understanding was a very useful process. Crucially, the fact that it was done as a team created a sense of ownership among all participants and resulted in consensus and deeper understanding of the whole process. Even once a high-level model was formed, participants understood the hidden structure that lay behind it.

The Model

This is the high level qualitative model that came out of the group's process. It involves the participation of different departments and represents the group's different mental models of how the market share of HP's product is affected by the price they set for that product, as well as by the overall market dynamics and consumer and retailer behaviour.

This model is too large for us to look at in detail here, but it will be interesting to look at a small part of it:

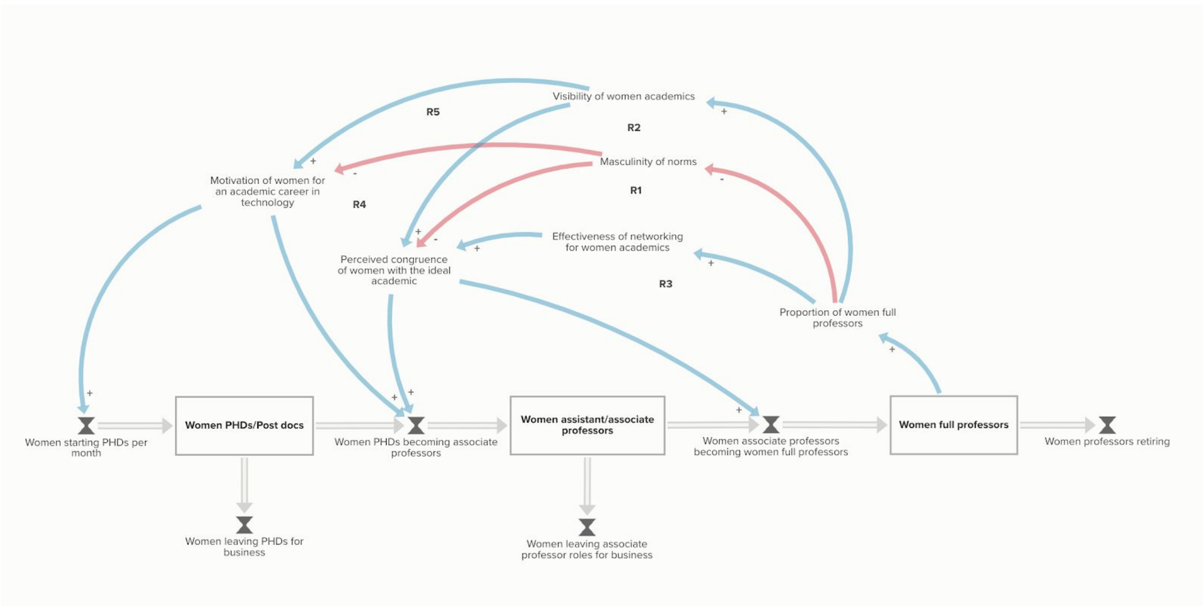

We can start by looking at the central variable in the model, which was the Probability purchaser buys HP supplies. In this case, the purchaser is most likely to be a retail store, who will then sell the product to consumers.

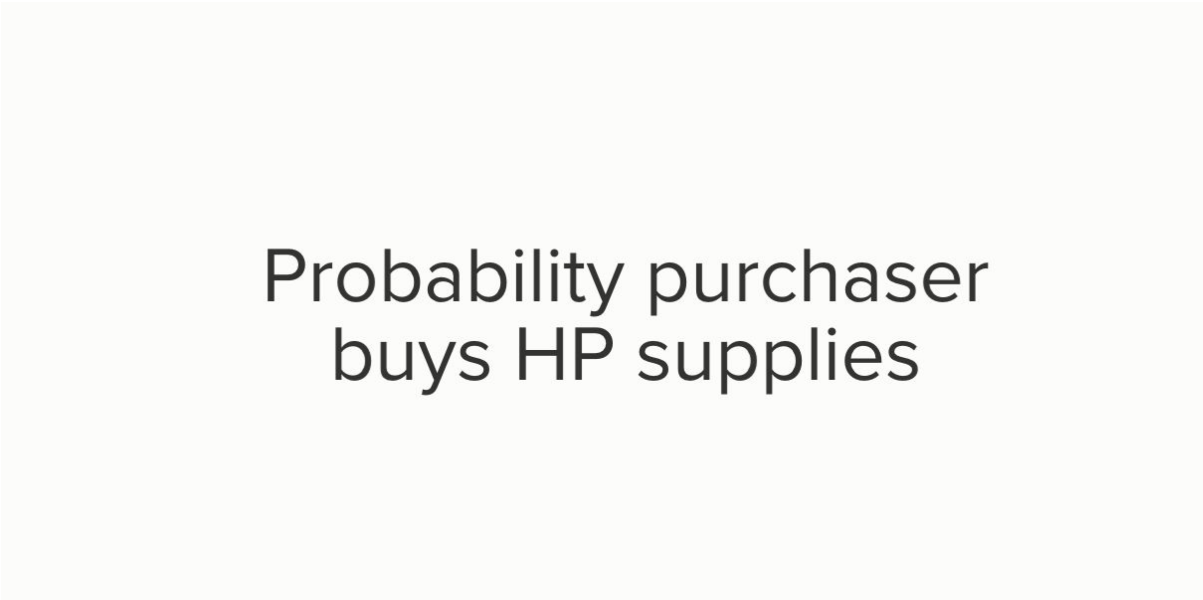

The team argued that the Probability purchaser buys HP supplies is affected by the End user and purchaser's satisfaction with HP supplies.

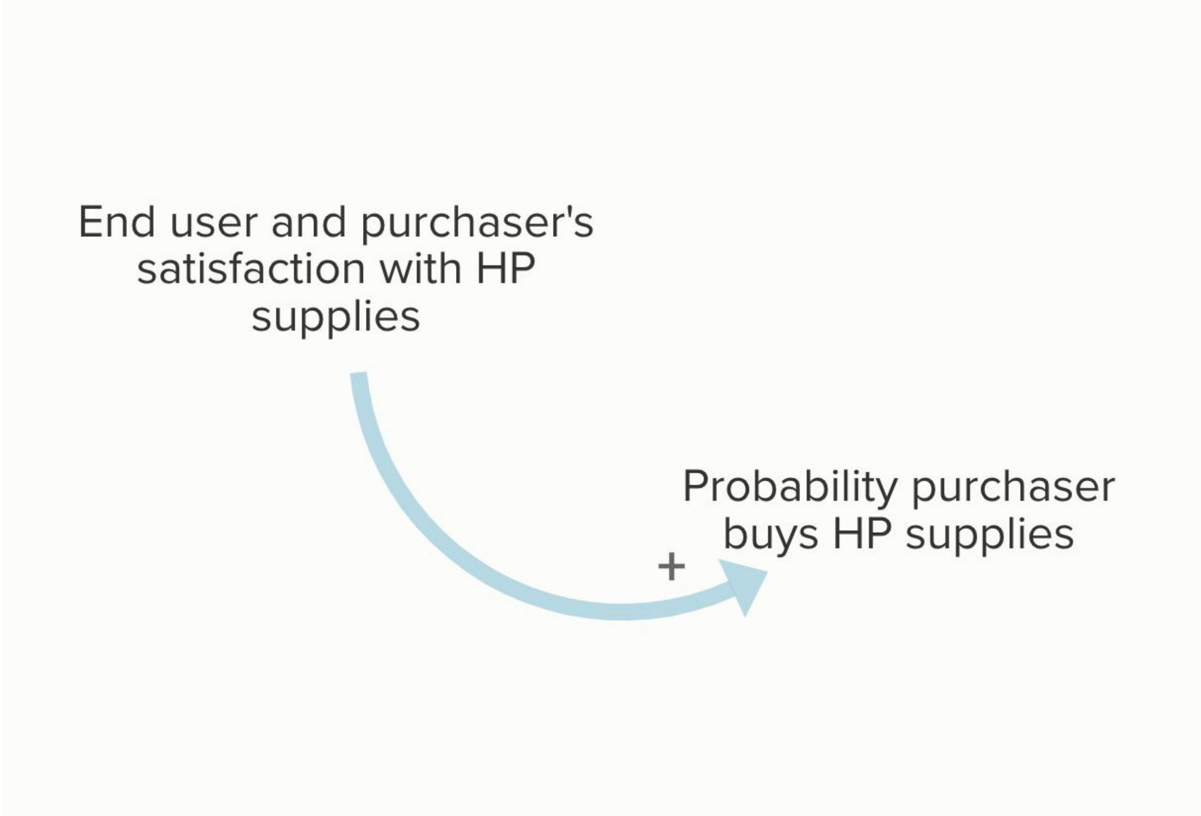

At the same time, the Probability purchaser buys HP supplies is reduced by their satisfaction with competitor's supplies. The more satisfied the retailer is with competitor's supplies, the less likely the retailer is to purchase supplies from HP. In essence, this is simply saying that HP are competing for market share with other firms.

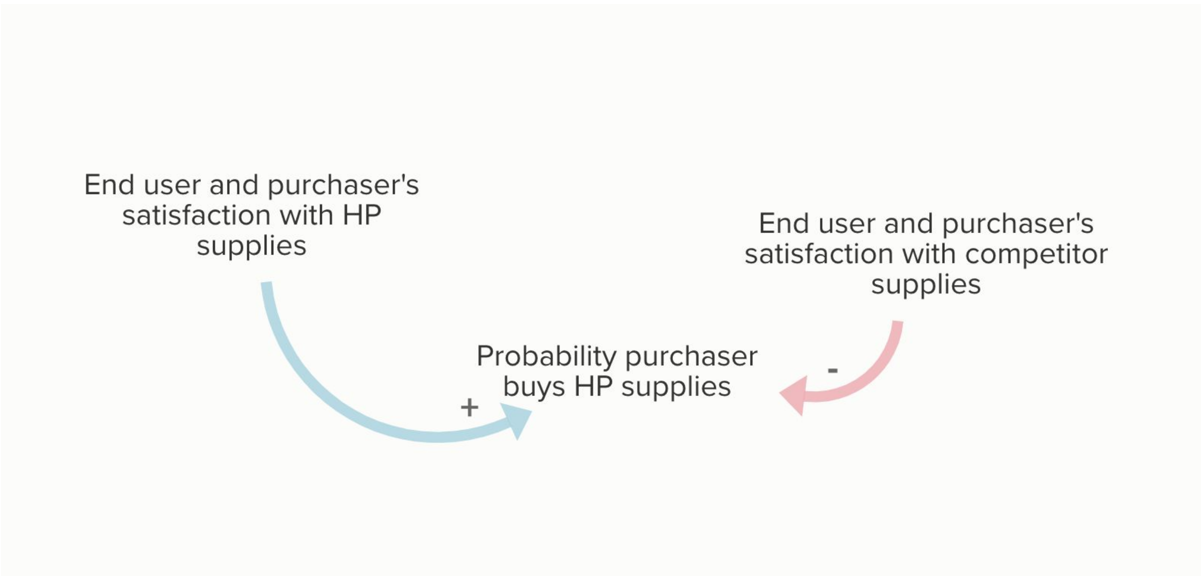

Ultimately, the probability of the purchaser buying HP supplies over their competitor's supplies will affect HP's sales, which will affect their profits, as modelled to the right.

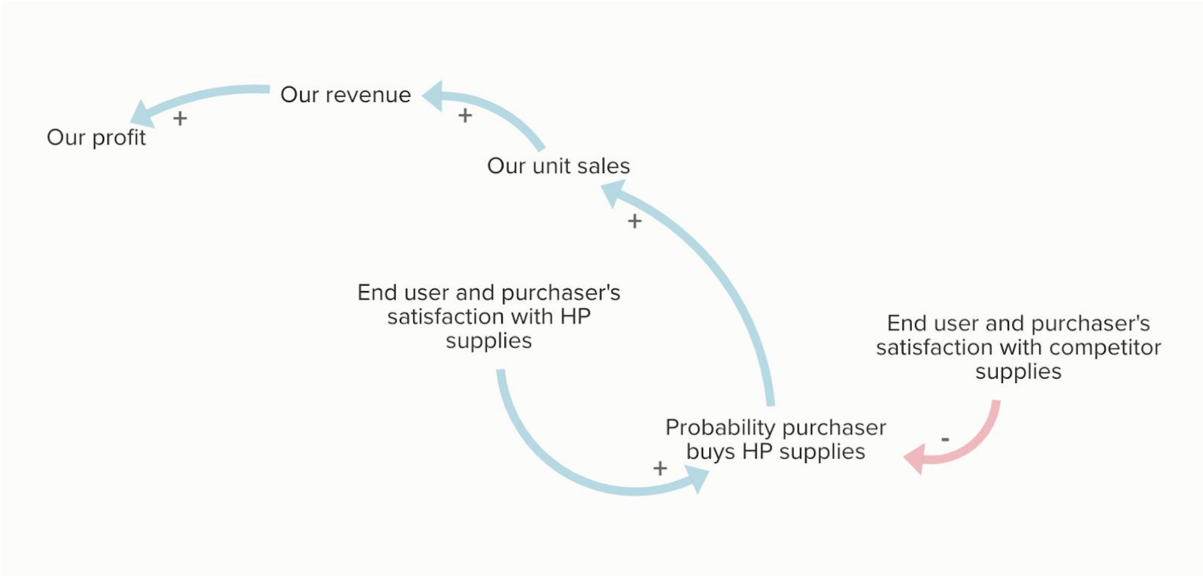

The team recognised that their profit will affect their investments. This is turn will affect their ability to communicate relative value to the retailer. In other words, an increase in profit will allow them to invest more in marketing and sales.

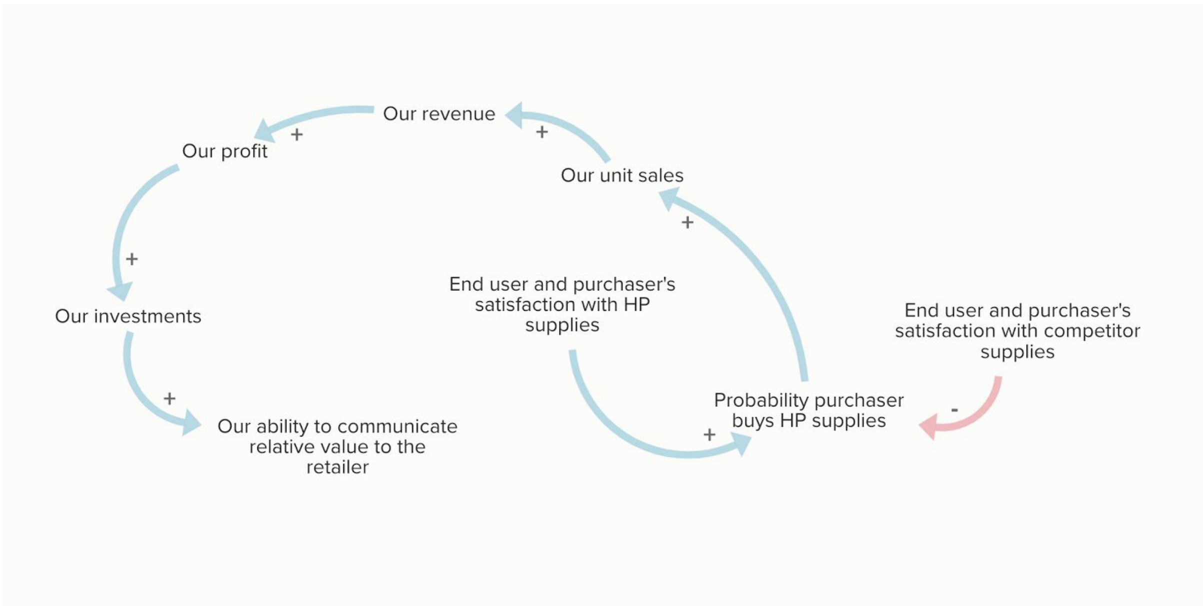

This in turn will increase the Retailer perception of our relative value proposition, which will increase the Probability purchaser buys HP supplies. This creates a reinforcing loop in the system, whereby HP's profits determine their ability to invest in marketing and sales, which in turn affects their sales and so their profits. When working as a virtuous loop, this will increase HP's market share exponentially. However, if competitor's get a better grip on market share, this will hurt HP's profits and so their ability to invest in marketing and sales, potentially sending their profits into a downward spiral.

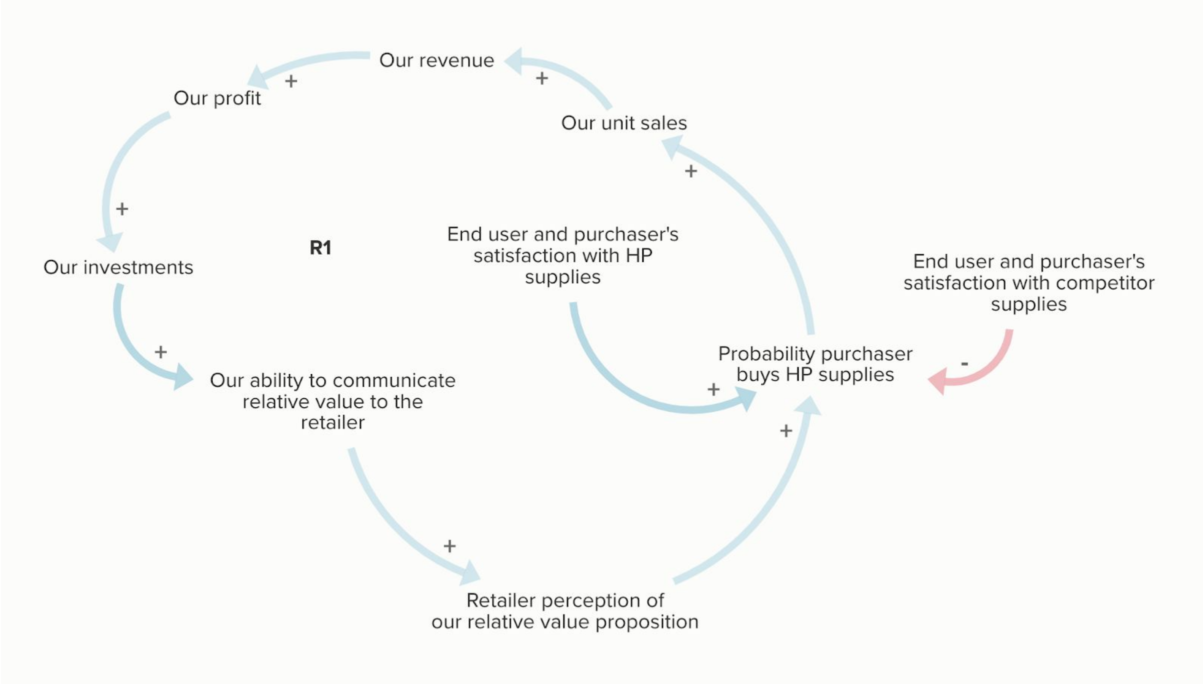

Outcome: Consensus and Action Plan

The management team was extremely impressed with the high-level model, and used it to develop a formal quantitative model for policy evaluation. Moreover, the model created a common communication tool among different departments to discuss different strategies and outcomes. The model was also helpful in bringing about many useful insights. For example, participants in the process explained how:

‘‘The model has highlighted the complexities, interdependencies, and gaps of knowledge about our business.’’

‘‘The factors enabling consumers to switch [to competitive supplies] are not being addressed aggressively enough.’’

‘‘We were reminded of the importance of influencing both supply product attributes and the consumer perception of supply product attributes simultaneously.’’
Moreover, the management found the model development journey to be useful. The journey through confusion and uncertainty created a shared understanding among all departments about how the overall system works, and how each department has a role in the company's market share. Furthermore, the model showed the interrelationships and complexity of the

overall structure, emphasizing the importance of a unified action plan. Crucially, developing the model with stakeholders created ownership and facilitated the implementation of the resulting action plan.

Many employees also felt that the model helped them to see the bigger picture of the business they were in. One commented about how the model ‘‘ makes you think beyond your job’’ and another spoke of how the model helped them to ‘‘see how what you do and do not do, affects the business’’.

Reflecting on Case Study 2

The HP case study demonstrates how CLDs can be used as a means of synthesising information and perspectives from different departments in a business. Interdepartmental communication is increasingly recognised as important for business success, and cases such as these demonstrate how CLDs can serve as a valuable methodology for structuring this communication.

To read more about this case, please see this article which describes the case.

Case Study 3: The Martina Bustos Community

The Context
The case of the systems modelling with the Martina Bustos community was described in one of the final sections in Module Four. In case you missed it there, the below text describes the context of the situation once more.

Martina Bustos is a settlement that is located 4 km from downtown Liberia, the capital of the Guanacaste region on Costa Rica. This settlement was established about 15-20 years ago when a private landowner, named Martina Bustos, donated the territory so that low-income families could have a place to live. However, Martina never formally handed over the land titles, and to this day they remain with her children.

Image courtesy of The Sustainability Laboratory

Nonetheless, many poor migrant families have moved to the area, and the population on the settlement has grown to between 3000-3500 people, mostly of Nicaraguan descent. The lack of land titles has caused many problems as, without them, the community is not formally recognised by the Costa Rican government. As a result, there is no legal obligation for the government to provide access to basic services such as electricity, waste management, and clean water.

In the summers of 2014 and 2015, participants in the Sustainability Lab's Global Sustainability Fellowship program began working with the Martina Bustos community. The purpose of the work was to create a community development plan that would serve as the blueprint for creating a more robust and sustainable community. The fellows worked with members of the community (mainly the women's association) to identify the numerous problems in the community, and try to find a strategy for resolving them. Much work had already been done by other NGOs in this area, but The Lab's approach was unique in that it took a holistic/systems perspective on the community, addressing multiple problems at once (economic, social, environmental, motivational) rather than focusing on just one issue at a time (for example, focusing on just housing or just water, as had been the focus of some previous work in the community).

Image courtesy of The Sustainability Laboratory

The use of CLDs

As part of this work, the fellows held Group Model Building sessions with the community members, the purpose of which was to discover together how the problems of the community might be interlinked, and how solutions to one issue would not suffice if considered in isolation to other issues. For these sessions, about ten community members were invited to the local university. They participated in a two hour session which was facilitated in Spanish by Vanessa Armendariz, one of the fellows.

Image courtesy of The Sustainability Laboratory

Vanessa began by asking the community members which issues in the communities were most important, which was decided on by a voting system. After a brief explanation of the purpose of the session, and a brief introduction to CLDs, Vanessa then wrote these issues on a white board, and began asking the community members to identify the immediate causes and effects of these issues. As they spoke, Vanessa began drawing arrows from causes to effects, stopping to ask participants if she was properly representing what they were saying.

A few hours of intense discussion later, the community had developed their own CLD that showed how the seemingly separate issues of the community were in fact heavily interlinked. The CLD also helped to identify many vicious cycles happening in the community, which could be turned into virtuous cycles if acted upon in the right way. New iterations of this CLD were then used by fellows as a basis for creating a strategy of community development in Martina Bustos.

The issues addressed in this session were quiet broad, with everything from the land rights to education levels to environmental problems (such as waste disposal) being represented. For this case study however, we will focus on just one part of the CLD, which represented the dynamics around community organisation and the land rights issue.

One of the issues identified by the community was the number of people with land rights. The land titles remained with the children of Martina Bustos, and this put the community members in a precarious position, as they felt they could be removed from their homes at any moment.

One effect of this was that the sense of belonging of each individual was very low.

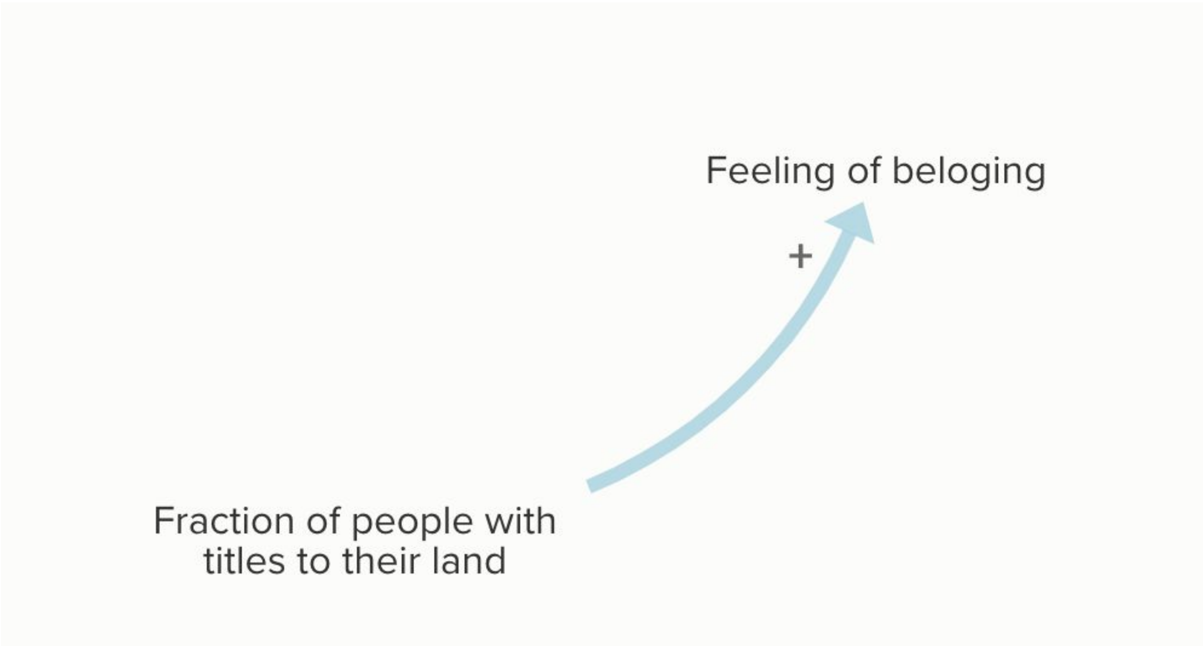

This lack of individual belonging then decreased the overall sense of community, as people often felt that their people were transient rather than permanent community members. As such they felt there was little sense in getting to know and rely on each other too much.

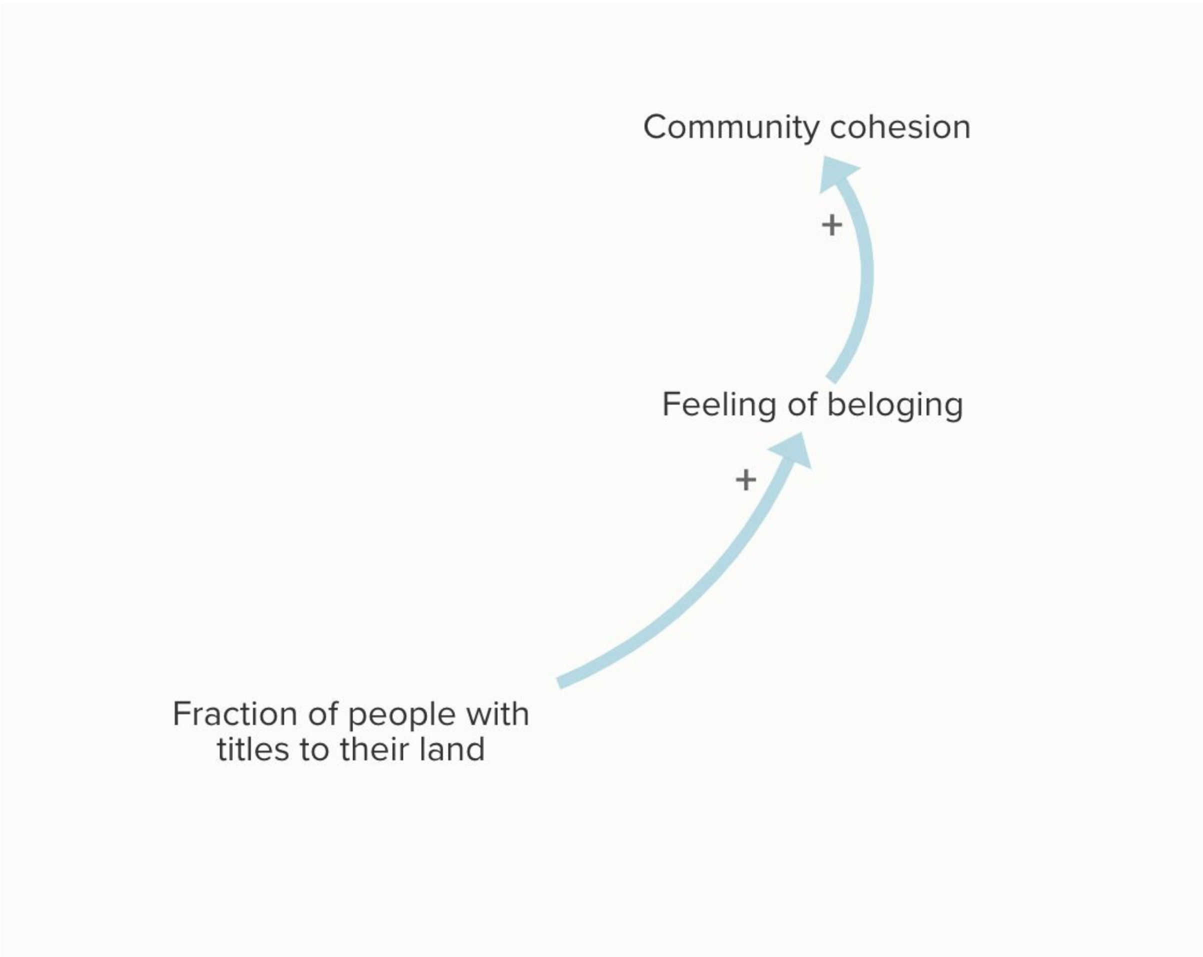

The lack of sense of community was a big factor in the lack of community organisation. For example, there was only one organised group in the community, known as the women's association, and they struggled to retain more than 5 or 6 regular attendees to their meetings. Many residents of the community said they were either not interested enough to attend these meeting, or were simply too busy. Moreover, many said that until they got their land titles, there was little point in putting effort into improving the community.

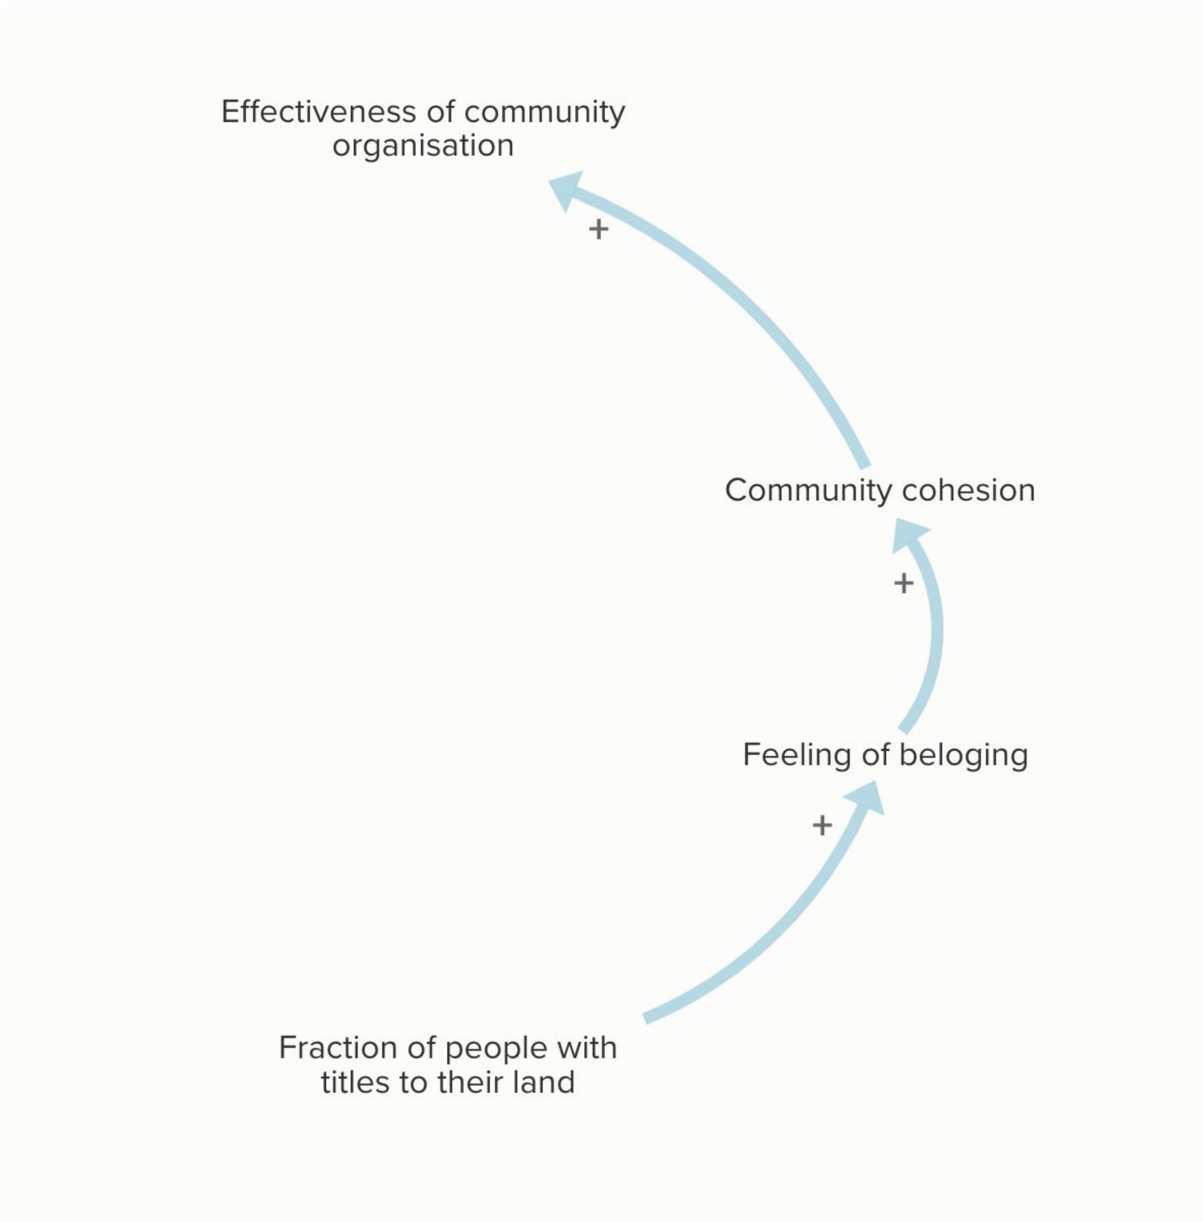

However, at the same time, many community members identified that the owner of the land titles, Martina's son, would be much more likely to hand over the land titles if he saw a more organised community. Indeed the son had often expressed doubt about whether handing over the land titles to such an unorganised community was a good idea.

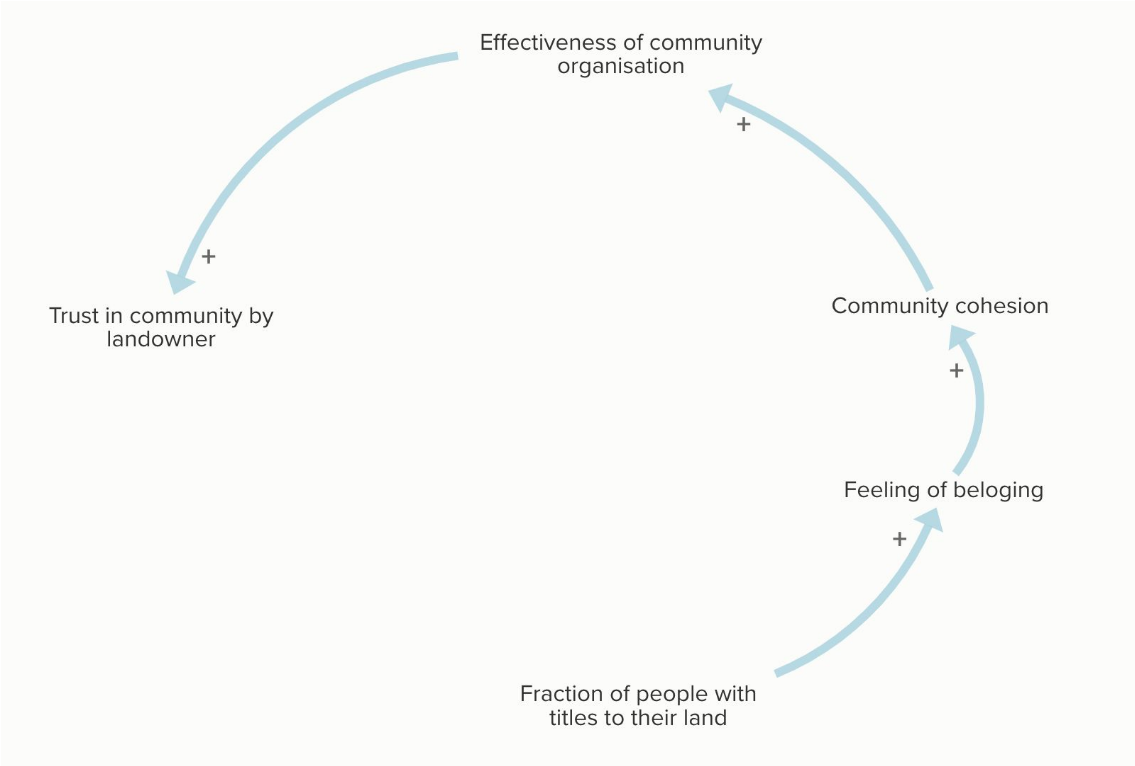

In identifying this, the group realised that there was a reinforcing feedback loop that was currently working against them. They saw how the community were not likely to organise themselves unless they got land titles, due to lack of communal sentiments. And at the same time, they saw how the land owner was not likely to hand over the titles until the community became more organised. Both the landowner and the community were stuck at an impasse, and one of them would have to move first to make things work

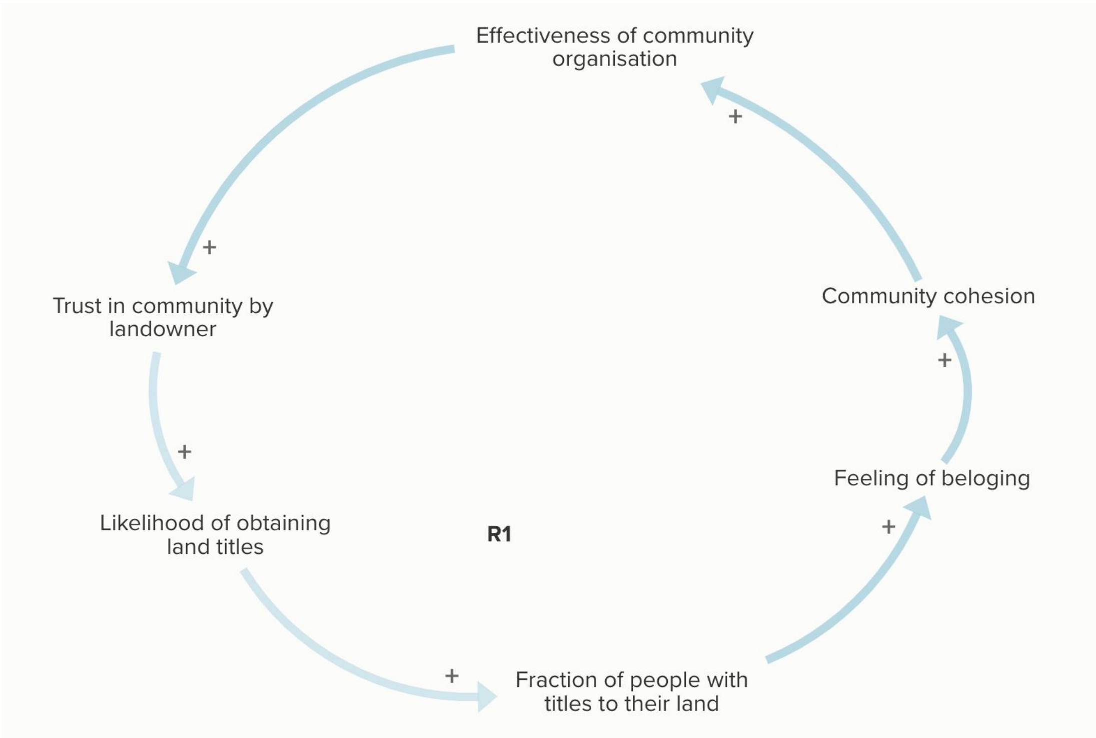

.

In addition to getting the land titles, the group also talked about how good community organisation can make them more likely to get government support for community projects. Currently, such projects are not supported by the government, because they don't believe the community can be organised enough to see their projects through and to maintain them when in place.

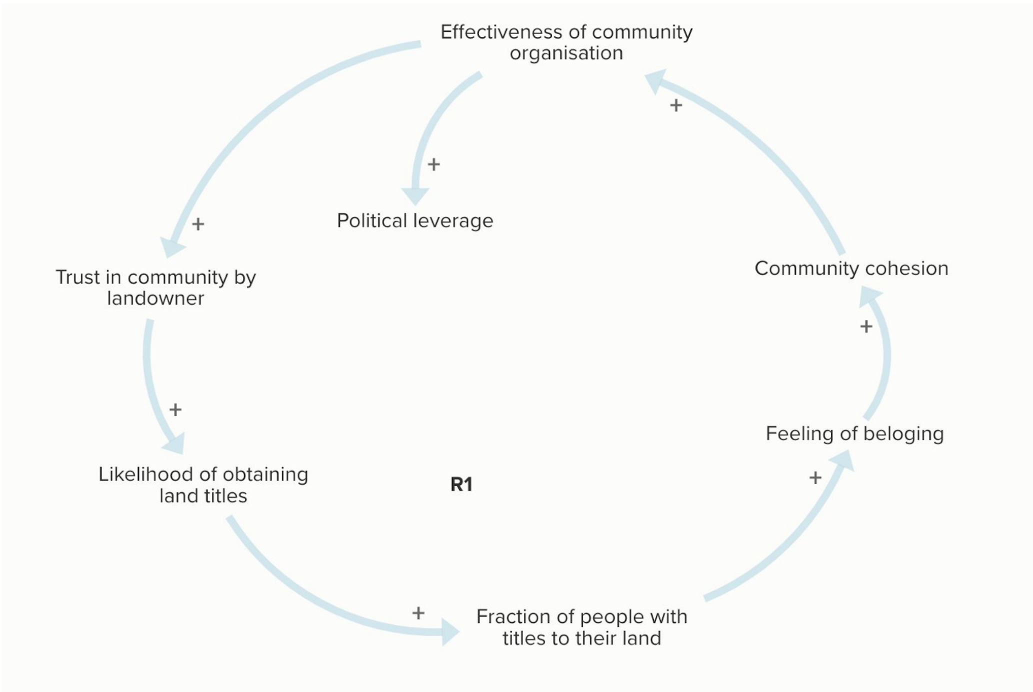

This lack of political leverage makes it less likely that there will be successful community projects, which in turn reduces the sense of community cohesion, further exacerbating the situation.

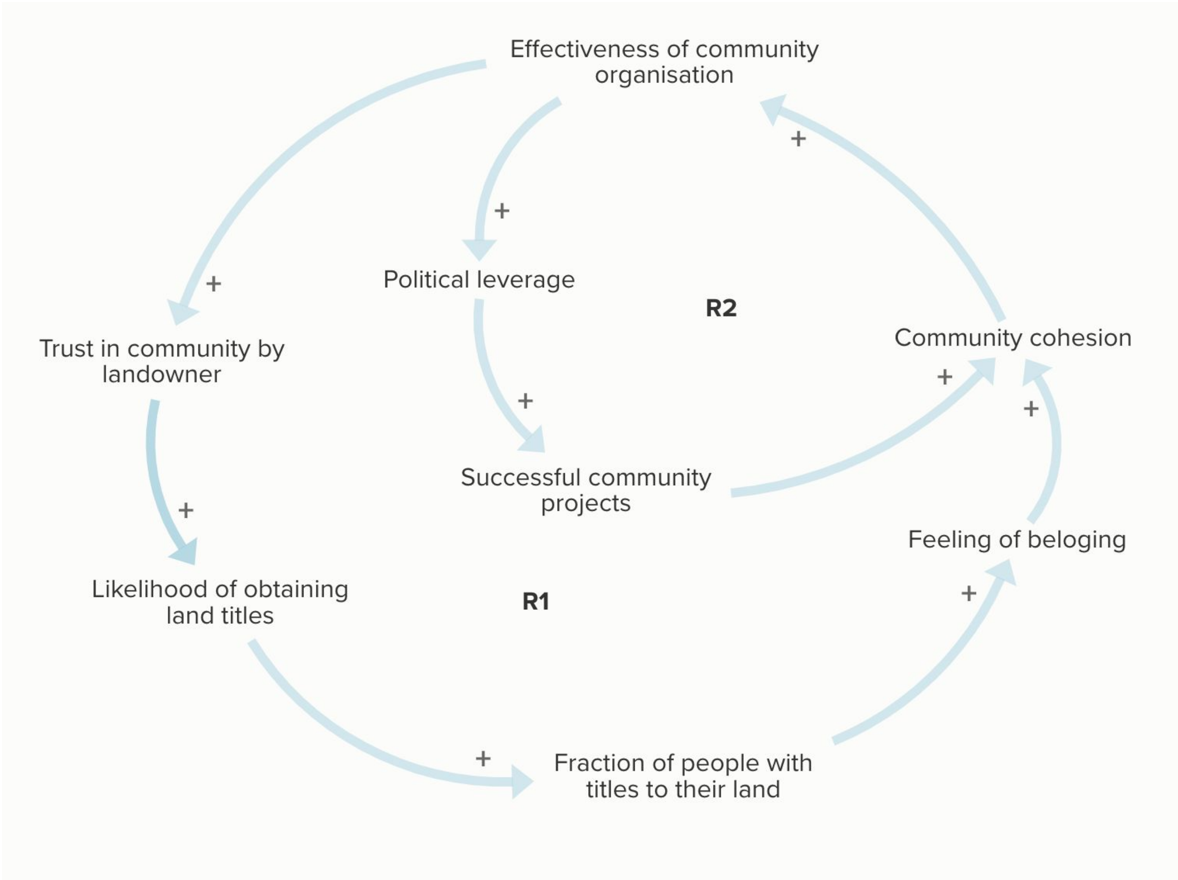

.

At the same time, the community pointed out that a lack of successful projects had led to a lack of community infrastructure. For example, there were only three or four places to get clean drinking water in the community, and the school was quite far away. Tasks like fetching water and getting the kids from a far away school were quiet time consuming, and this reduced the time they had for attending community meetings and helping out with community projects, which was another block towards communal action.

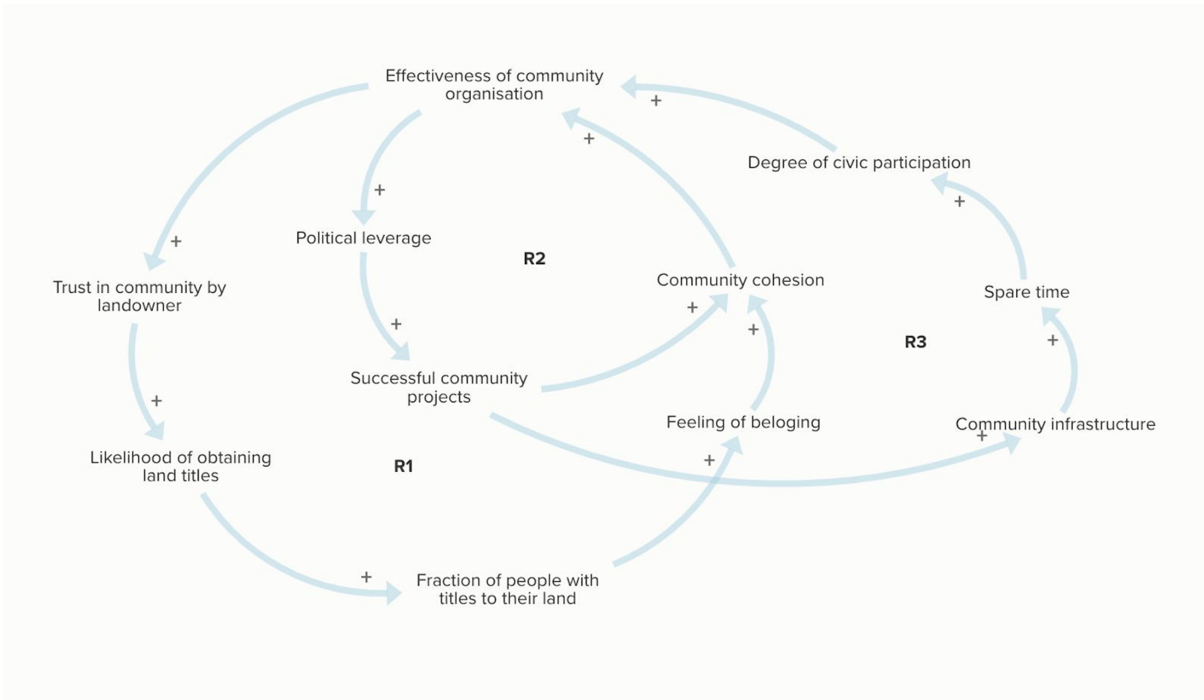

In addition to a lack of successful community projects, the lack of land titles meant that the government had no mandate to provide access to basic services such as clean water or electricity.

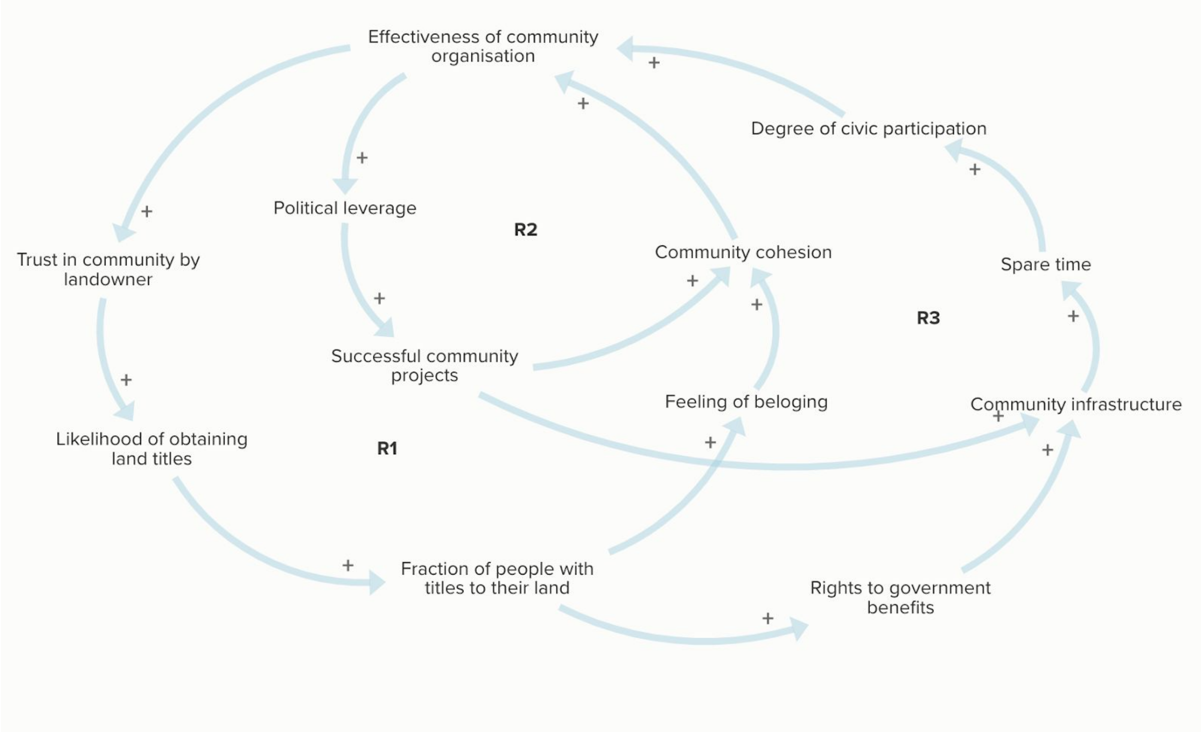

With the identification of all these interlinked feedback loops working against them, the community could see an opportunity to turn things around. They could see how more community participation could lead to more likelihood of receiving the land titles, which could increase the community participation of others. At the same time, they could see the potential for political leverage, which in turn could help with community infrastructure. Infrastructure could also help by freeing up community members' time that would otherwise be spent fetching water or collecting their children from a far away school.

Outcomes

One of the most important outcome of the use of CLDs in this process was the dialog that it facilitated between community members and the fellows. Fellows were able to dig deeper into the communities issues, and community members were able to hold a fruitful discussion about the inter-linked nature of their issues. Simply having a facilitated conversation with the common goal of producing a model of the community's problems seemed to allow for a more fruitful discussion than usual, which is indeed one of the strengths of group model building.

Additionally, the CLDs played a crucial role in the formation of the strategy that was created by and with the community. This strategy has been described in this document, which outlines a blueprint for improving conditions in the Martina Bustos community. The Lab is currently seeking political and financial support to help implement this strategy.

Reflection on Case Study 3

The Martina Bustos case study demonstrates how causal loop diagramming can help a community to not only communicate more effectively between themselves, but also to synthesise their different perspectives on an issue so as to realise new solutions. In this case, the community realised the chick and egg situation they were in in regards to the land titles issue. They saw how the land owner would not hand over the titles until the community organised itself more, but they also saw how many community members were not willing to participate in community building activities until they had the security of the land titles. Seeing this as a reinforcing feedback loop was an insightful moment for the community members and for the fellows of The Lab.

Case Study 4: The ISLANDS Project

Introduction

Small island states are specifically vulnerable to climate change, natural disasters, and external shocks such as energy crises and financial volatility. Part of this is due to their small size, lack of access to markets, financial resources, and their limited human and institutional capacity.

This vulnerability has been specifically recognized in various environmental and development strategies. As a result, initiatives are forming to increase the resilience and achieve sustainability objectives in these states. One such example is the ISLANDS project, aiming to contribute to an increased level of social, economic, and ecological sustainability, while simultaneously strengthening cooperation, in the region of Southern and Eastern Africa and Western Indian Ocean. More specifically, the project was carried out in Comoros, Madagascar, Mauritius, Zanzibar, and Seychelles. Here, we will go through the process, have a look at how system dynamics was used, and explore the results and outcomes of the process for Seychelles. For the full story, see this article, by Deenapanray and Bassi (2014), or section 4.3.1 of this book by Gilbert Probst and Andrea Bassi.

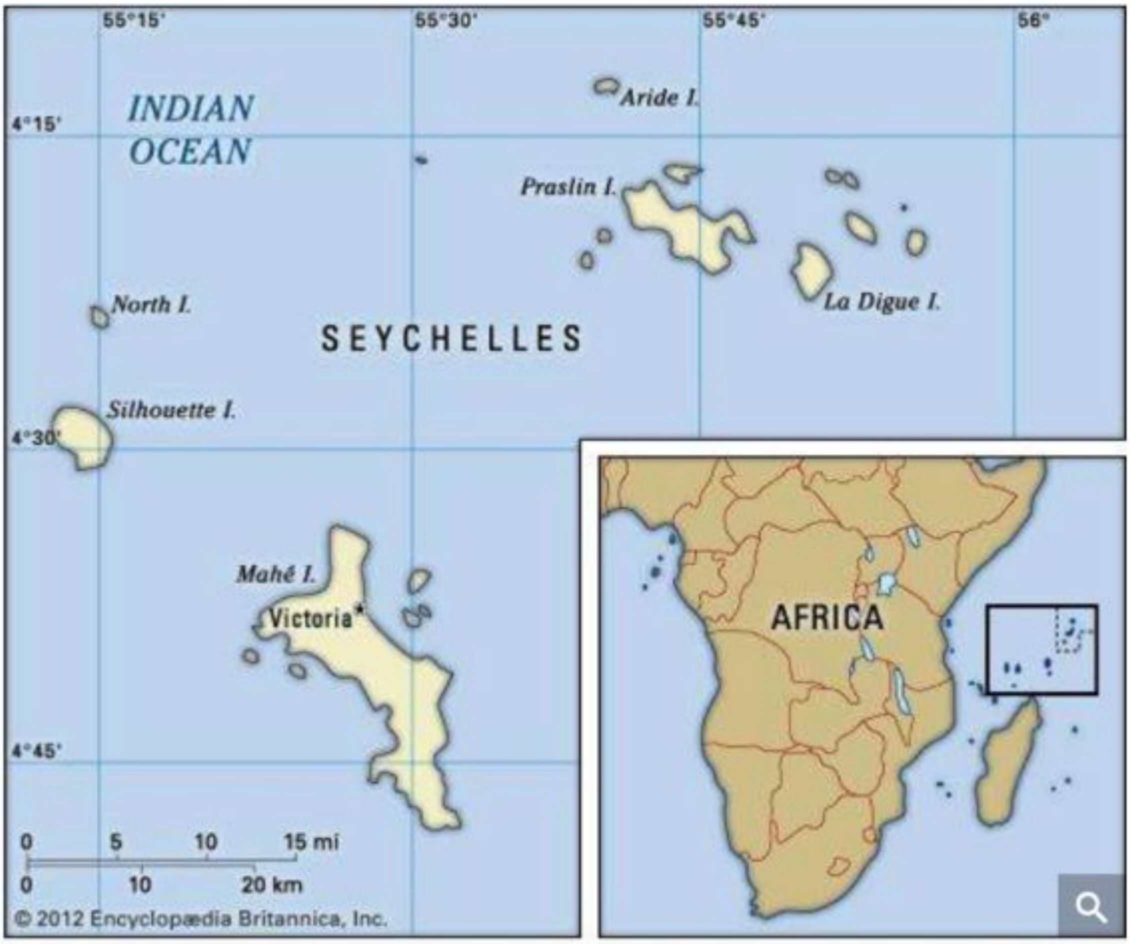

Figure 1. Seychelles is one of the world's smallest countries, compromising 115 islands, and located in the Indian Ocean - east of Kenya and north east of Madagascar. The country holds a remarkable biodiversity, where the wildlife includes sea turtles, tree frogs, chameleons, and over 900 species of fish. Image courtesy of Encyclopaedia Britannica.

The Project Design

The ISLANDS project was designed based on three pillars. The first important underpinning was resilience and complex systems theory. When employing this perspective, sustainability is seen as an emerging property of complex social-ecological systems, where the views and conceptual understanding of sustainability among the actors in the systems shape the future development. Moreover, the project made use of a multi-stakeholder process, to establish cooperation and a community of practice among the participating countries. Lastly, capacity building and learning-by-doing were crucial to the project development.

The ISLANDS project consisted of four key steps. First, a monitoring system was set in place, to evaluate the implementation of the sustainability strategy on national, regional, and international levels. Second, best practices in mitigation and adaptation to the vulnerabilities of the island states were identified through data collection and modelling. Policy options and potential interventions were assessed using system dynamics modelling, through a lens of local culture and values. Therefore, it was specifically important to acknowledge different disciplinary perspectives, and to include both expert and non-expert knowledge in the process. Third, a commitment to action was established, and financial support leveraged. Lastly, the partnership between the actors engaged in the project was strengthened.

The Modelling Process

The aim of the modelling process was to build ownership and trust, shared knowledge about the policy context in the different countries, and from there conduct policy analysis to support small island states in their work to achieve sustainability objectives. The participants in the model building sessions included representatives from academia, the government, civil society, and the private sector.

In the working group in Seychelles, there was in total 15 people taking part in the process. The CLD created by the group included the interactions between three economic sectors identified as important in the country; the fishery sector, tourism, and the financial sector.

The Model

The figure to the right displays a simplified version of the CLD, highlighting some key feedback loops, which we will discuss individually.

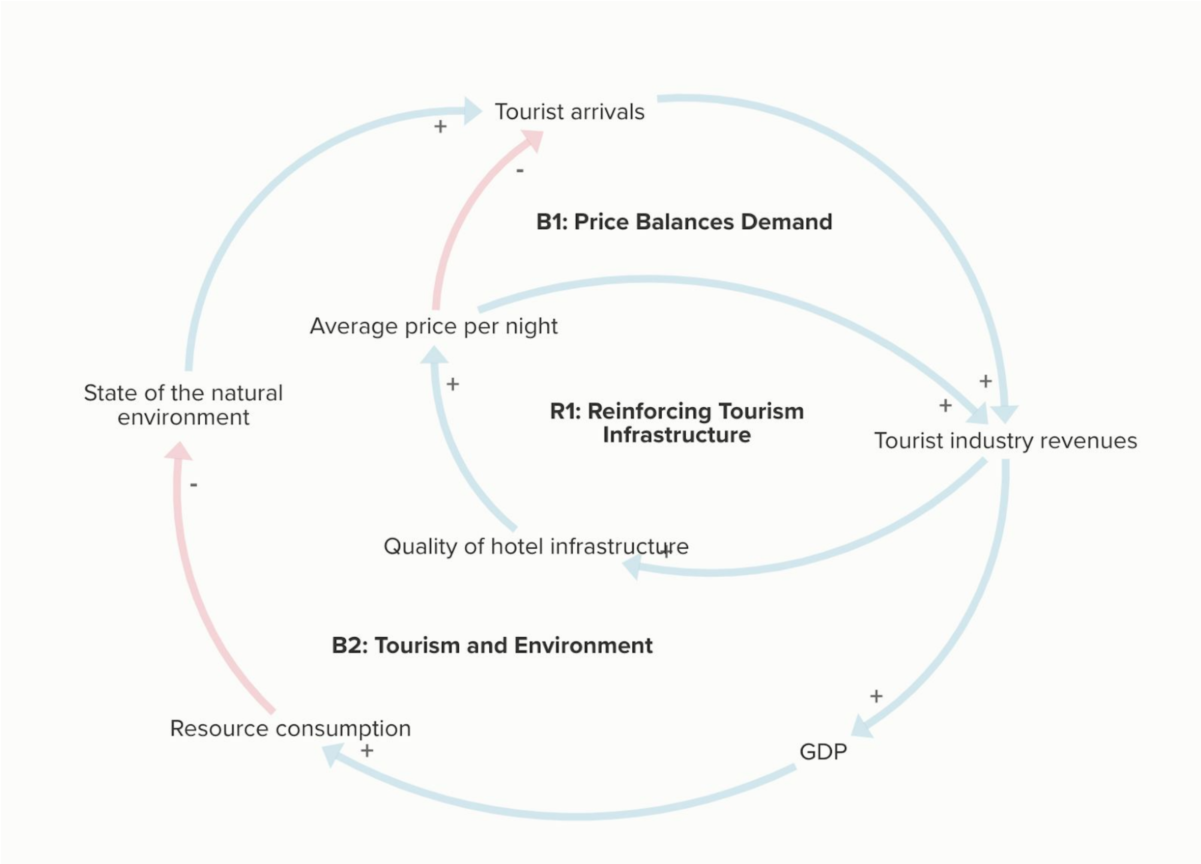

R1: Reinforcing Tourism Infrastructure

As depicted by the reinforcing feedback R1, an expansion of the tourism sector can be reinforced by making investments to improve service provision. The larger the tourist industry revenues, the greater the ability to invest in the quality of hotels and tourist infrastructure. This provides an opportunity to charge a higher price for hotel stays, in turn generating even larger revenues for the tourism sector.

This loop could also work the other way, whereby a decrease in Tourist industry revenues could reduce the Quality of hotel infrastructure, thereby reducing the Average price per night and further reducing the Tourist industry revenues.

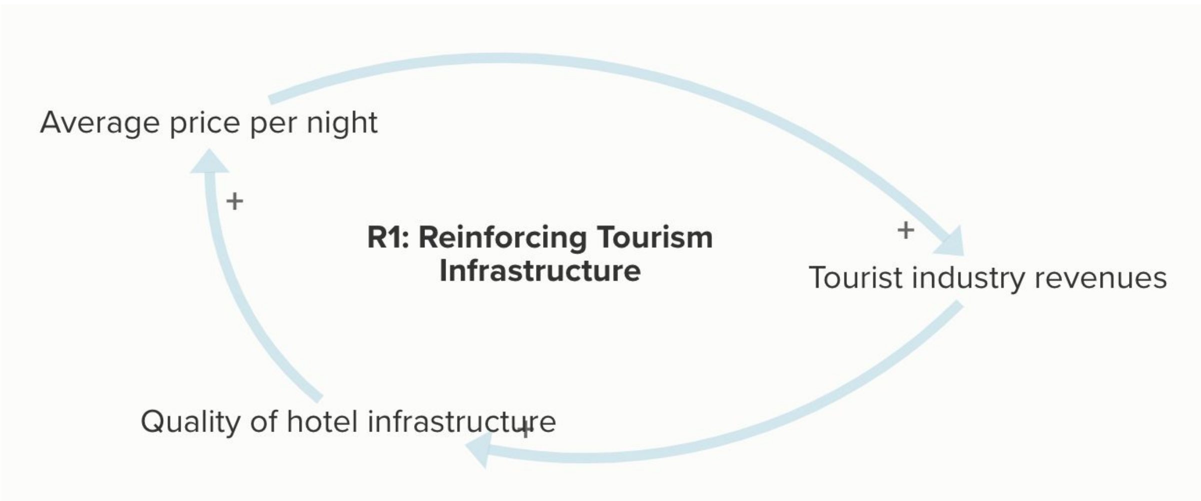

B1: Price Balances Demand

However, higher hotel prices will also deter tourists, thereby lowering Tourist industry revenues, thereby lowering the ability to make further investments in Quality of hotel infrastructure, which forces companies to lower the Average price per night. This is shown by the balancing feedback B2.

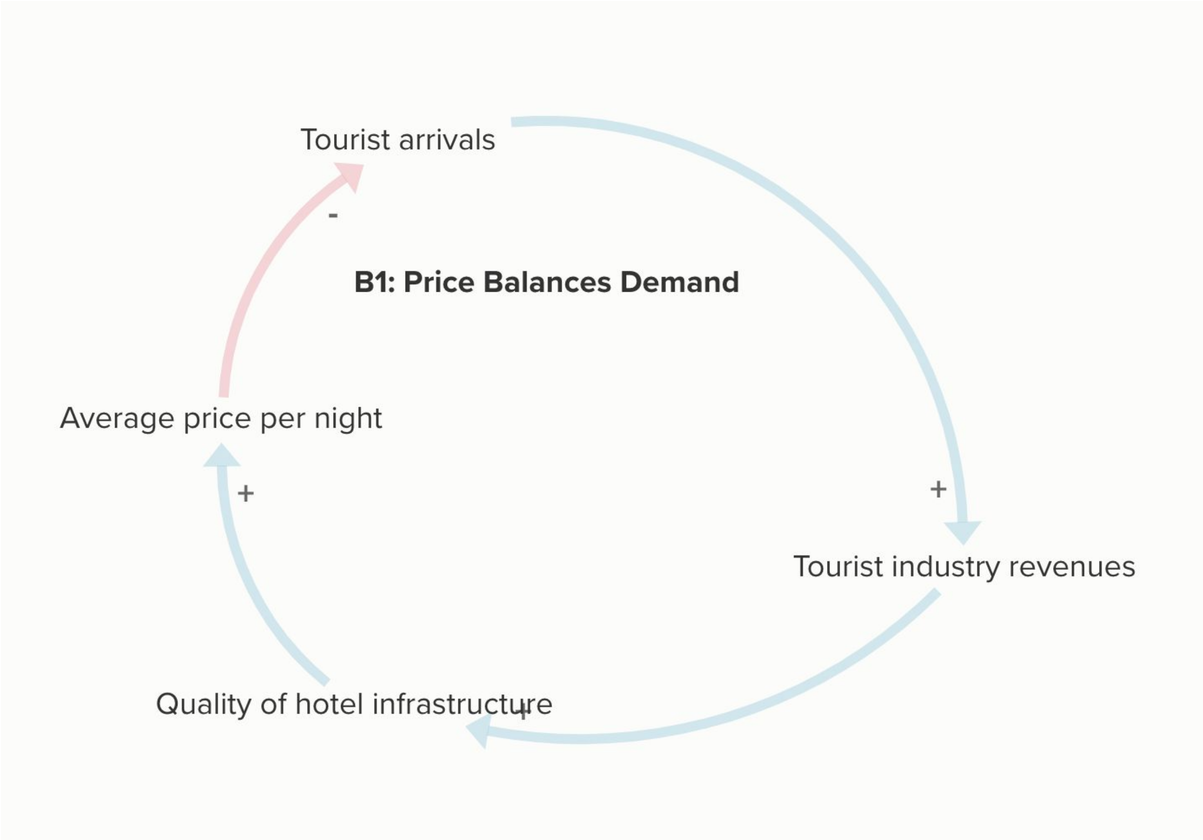

B2: Tourism and Environment

A main driver of tourism in Seychelles is the natural environment, offering beautiful surroundings and a rich wildlife. However, an expanding tourist industry results in increasing resource consumption and negative environmental impacts, thereby lowering the attractiveness of Seychelles as a tourist destination. This in turn reduces the arrival of tourists. This dynamic is represented by the balancing feedback B2.

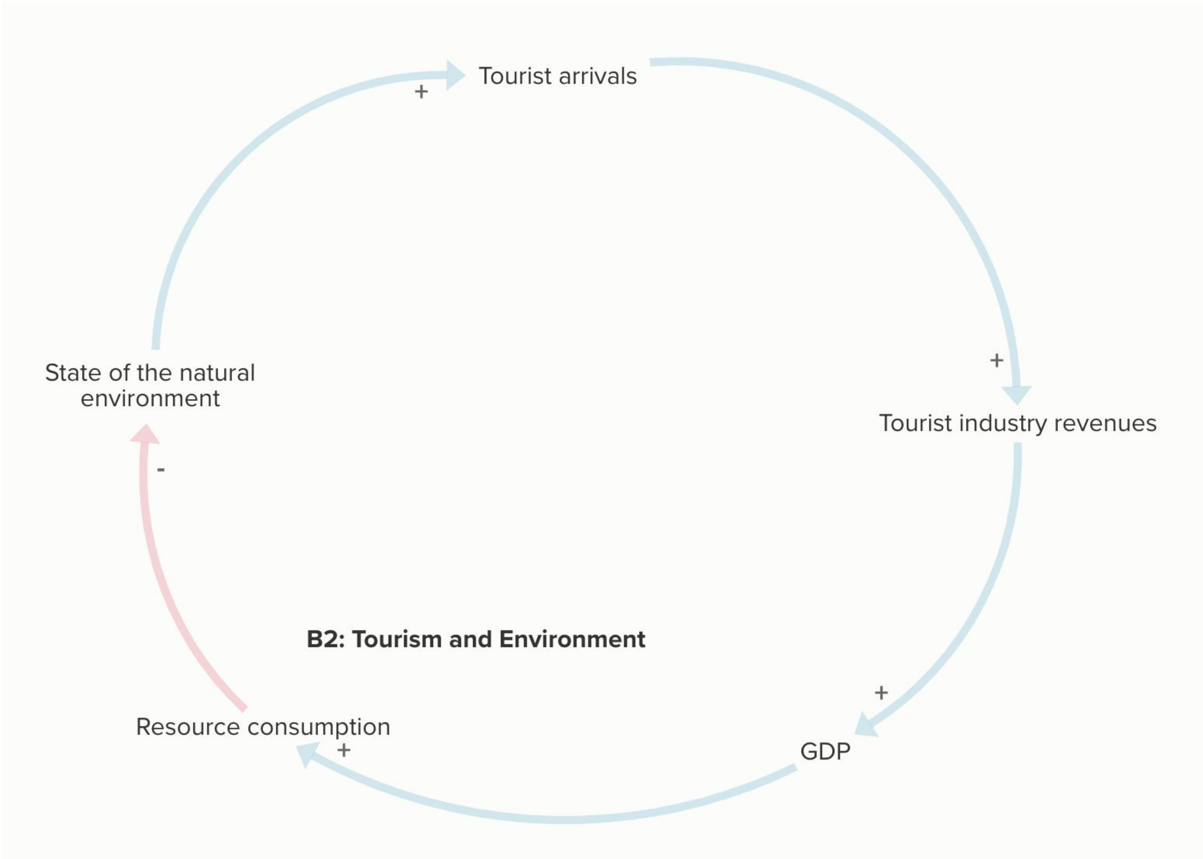

Putting these insights together we can see how there is a complex interplay between the size and quality of the tourism industry, and the state of the natural environment. In this case the CLD was much larger and represented many more sectors of the economy, but we just wanted to give a small preview of the kind of modelling that can be done in such cases.

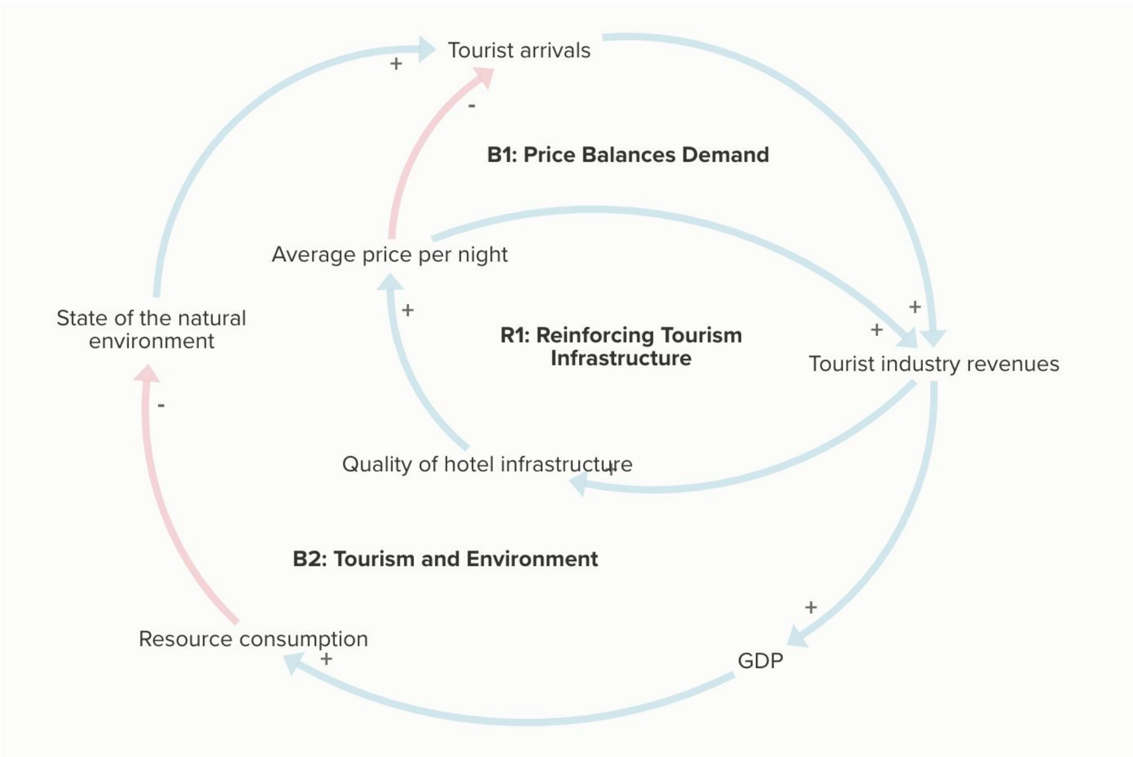

Outcomes

The ISLANDS project had multiple outcomes. First, new forms of co-operation were established, with expert groups from the participating countries now meeting regularly and engaging in knowledge sharing. Additionally, through the modelling process data gaps were identified, which could guide future research efforts. Lastly, steps were taken to institutionalize the use of system dynamics modelling as a tool to support policy design and testing, particularly within the context of national development planning for Seychelles and Mauritius.

Figure 2. The economy in Seychelles is heavily dependent on the tourist industry, representing around 25% of the total GDP. Image courtesy of Encyclopaedia Britannica.

Reflection on Case Study 4

The ISLANDS project demonstrates how CLDs can be used to bring various stakeholders together to analyse a problem from a multi-disciplinary perspective. Such analyses are increasingly recognised as vital to the sustainability movement, and the ISLANDS project offers another example of how CLDs can be used as a tool to help foster these more holistic understandings of complex social, economic, and environmental issues.

Case Study 5: Building a Healthcare System

Story Background

Healthcare systems are most certainly complex, and managing such a system is challenging. The task becomes even more difficult when one is faced with a decline in budget. This was the situation faced by Dr. John Hunter (JH). John was the lead manager and developer of an oncological care unit (a unit for treating cancer) situated in the department of internal medicine at the central hospital in Klagenfurt, Austria. He had pioneered and directed the development of this oncological care system over roughly 25 years , covering the whole state and involving ten hospitals as well as multiple local physicians.

In 2009, John learned that the budget for the oncology unit was going to be reduced, as the administration of the central hospital argued that it needed to "improve the economic situation."

The Methods and Frameworks Used

John had been using the help of an organizational scientist, Mark Stone (MS), since he first took up the challenge of developing the oncological care unit. They had worked in conjunction to conceptualize and organize the unit, and they had relied heavily upon the method of system dynamics and the viable systems theory (mentioned in Module One).

In this case study, we will focus on the use of causal loop diagramming - a method of qualitative system dynamics - in the project. However, if you are interested in learning about the wider theoretical context behind the study, then you may want to read this paper.

The Use of Causal Loop Diagramming

John had used system dynamics modelling from the outset of this project , as a method for visualizing and accounting for the interdependencies between the many different parts of the healthcare system, some of which were quite tangible (e.g. available infrastructure) and some of which referred to less tangible but equally important variables (e.g. staff motivation and the reputation of the oncology unit).

When the budget cut was announced in 2009, John decided that system dynamics could be a useful tool in analyzing the potential consequences. The results of this analysis could then be presented to the administration team at the central hospital, so that they would be aware of the implications of their decision.

The Model

Through their analysis and conversations with oncologists, JH and MS identified some noteworthy dynamics that would likely arise as a result of the budget cut. We can now explore this model one loop at a time.

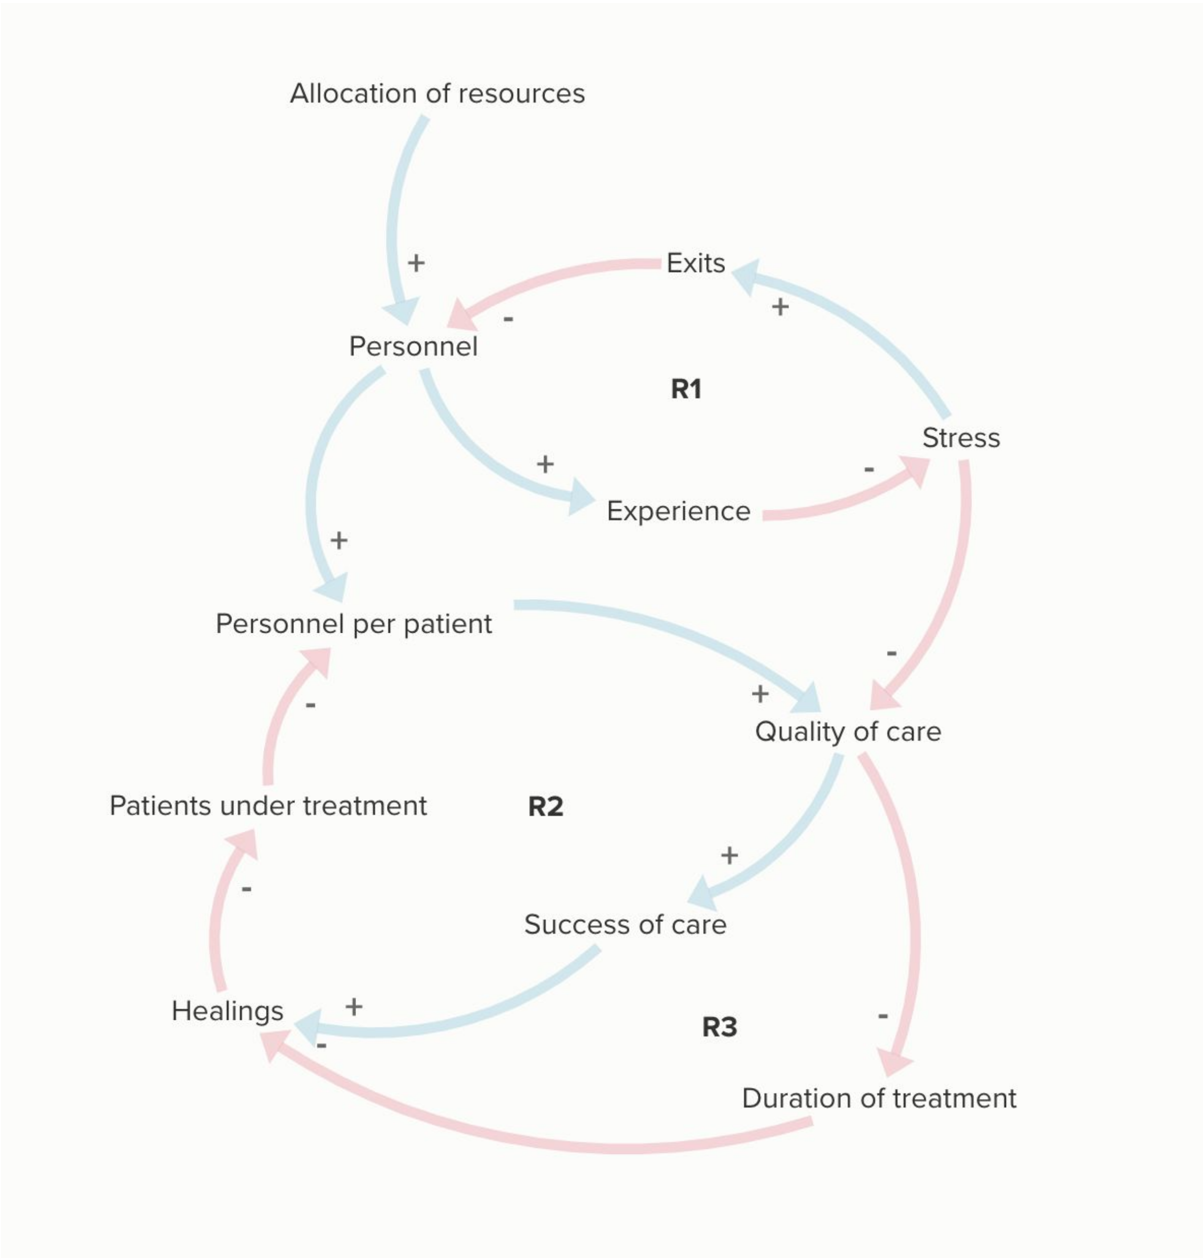

It's clear that a lower budget would reduce the Personnel (i.e., the staff in the care unit) available. This would then reduce the level of Experience available in the organisation, which would cause more Stress in an already stressful environment, as the staff's capacity to deal with problems effectively drops. The increase in Stress would then cause an increase in Exits, as overly burdened staff looked for new jobs, or have more stress-related absences.

This structure highlights the potential for a downward spiral to occur, whereby a small initial reduction in staff (due to the budget cut) could result in a continuous and accelerating decline in staff, as stress levels pushes more and more people out the door.

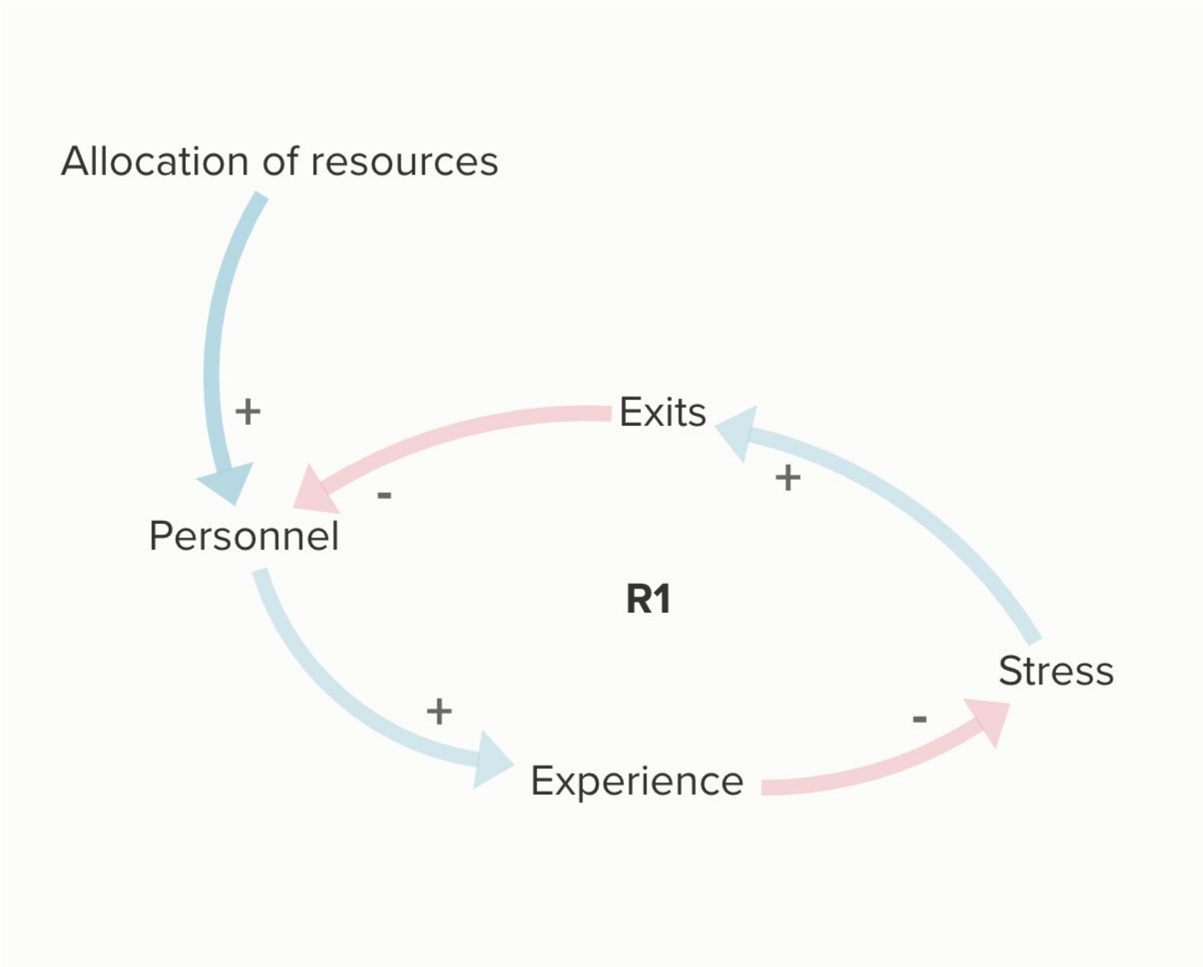

At the same time, a reduction in Personnel would reduce the patient-doctor ratio, which would reduce the Quality of care and thus the Success of care. This, in turn, could reduce the number of Healings, which would slow down the outflow of patients, and thereby further reduce the doctor-patient ratio, i.e. the Personnel per patient.

This highlights how the reduction in budget could cause another downward spiral in the system.

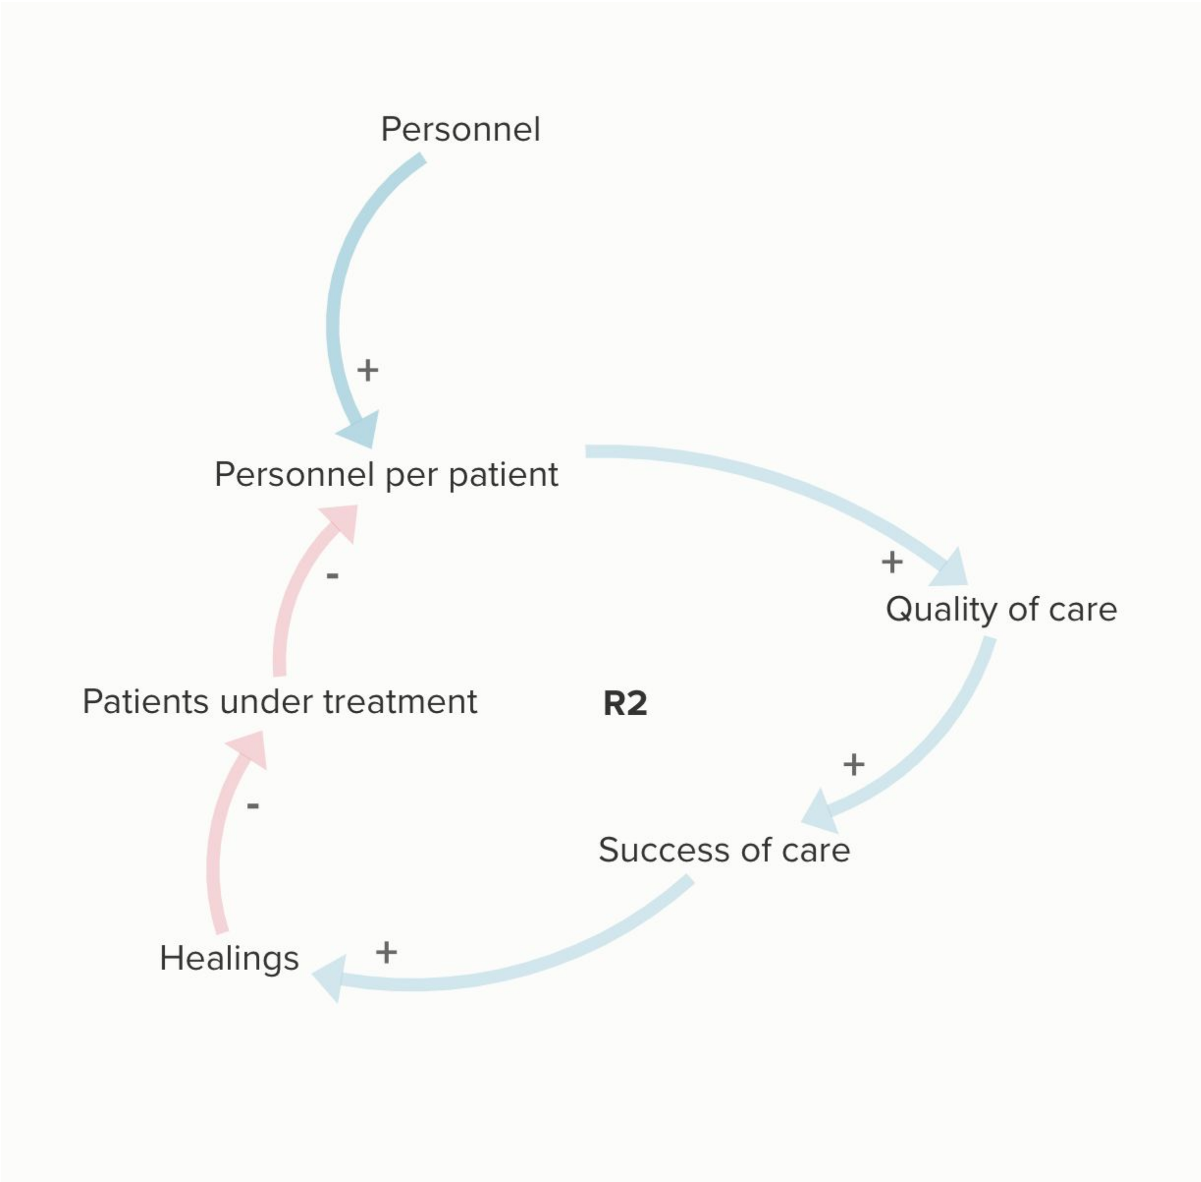

Similarly, a reduction in the Quality of care would not only reduce the Success of care, but would also increase the Duration of treatments. This would also slow down the Healings, causing the number of patients to accumulate.

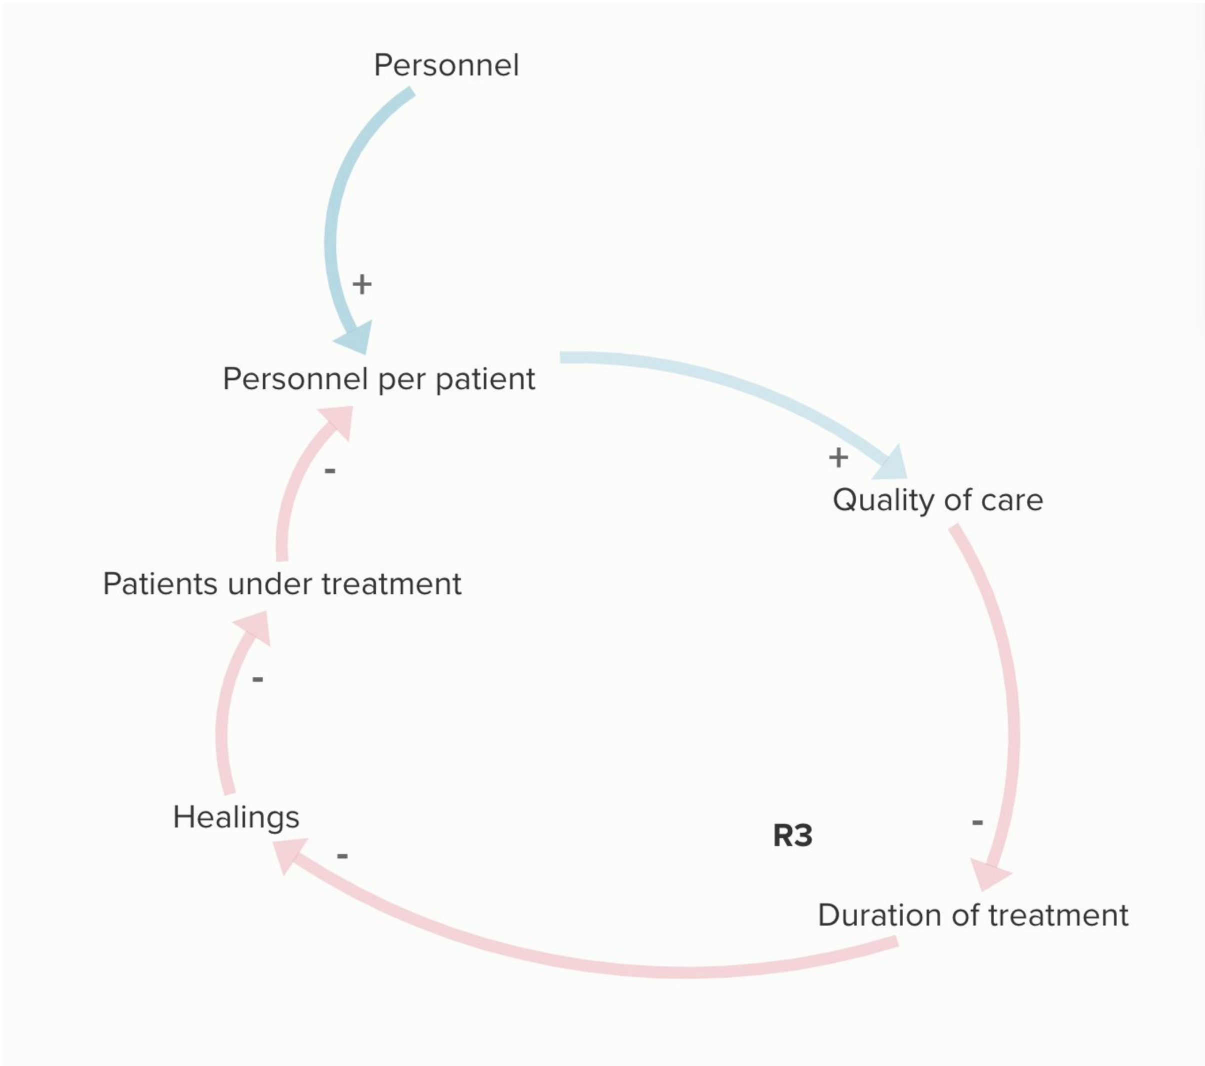

Taking a look at the whole model, we can see how loop R1 is also linked with loops R2 and R3, as the reduction in Personnel, and the corresponding increase in Stress, will affect the Quality of care in various ways.

Together, these loops represent the dynamic hypothesis of what the oncologists thought would happen as a result of the budget cut: a loss of Personnel and therefore Experience will be further reinforced due to the increasing overload and Stress and therefore Exits of the personnel in the ontological care unit. This will lead to a decay in the Quality of care, which will result in a growing load of patients to be treated. This not only further increases the burden on existing staff, but also increases the hospital's costs, as patients spend longer and longer in the hospital.

The result is that the budget cut would actually worsen the economic situation by increasing the hospital's costs, as well as reduce its income by slowing down the rate of successful treatments. Thus the model shows how the decision to cut the budget would have the opposite effect than desired, due to the dynamics described above.

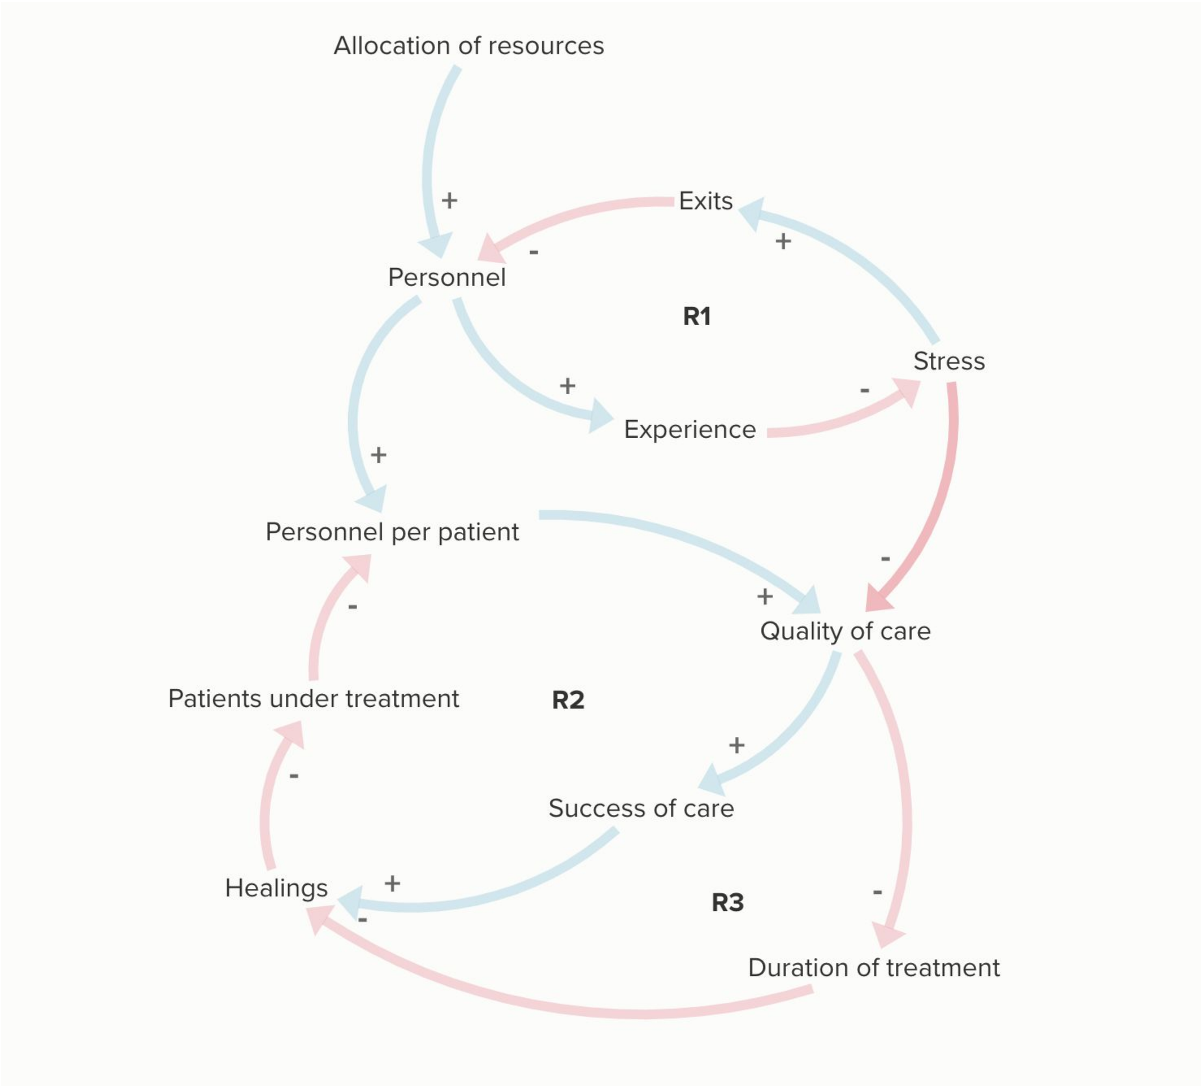

Outcomes

To provide a more rigorous perspective, a simulation model was built as a next step in the process. The results of this and the qualitative CLDs were then presented to the staff of the oncological unit, who agreed with the analysis and appreciated the insights gained from it. However, when the model and its insights were presented to the central hospital administration, it seemed to have little to no impact on their decision regarding the budget cut.

The authors of the case report speculate that there could be two reasons for this. One is that the administration teams were unreasonable and stubborn. The other, perhaps more elaborate reason, is worth hearing in full, as it represents an important piece of wisdom that has often been learned the hard way in system dynamics studies so far. The authors explained that:

"people are motivated by a sense of “psychological ownership,” a feeling that they can lay claim to certain organizational factors as their own (Pierce et al. 2001). Such ownership fosters self-identification as well as organizational commitment and citizenship behavior (Van Dyne and Pierce 2004). The managers who initiated the budget cut had in no way been involved in the formation of the simulation model. That model had been developed rather spontaneously between the two authors. They (the authors) did not consider, at that point, what in hindsight appears as a straightforward way to proceed: trying to involve the administrators in the construction of the model and using the model in the political negotiation process."

This led the authors to conclude that "the improvement of a management system can be best obtained not only by designing a dynamic, multidimensional model but also giving decision-makers a sense of ownership of the system under development through their participation in the design process. Good models and logically cogent argumentation, based on careful studies at the content level, are necessary to trigger change in direction of a desirable future. But they are not sufficient. A context needs to be created that makes these arguments effective. In our case, the management of the hospital was not prepared adequately. They enacted their decision without an adequate understanding of its implications."

Reflection on Case Study 5

This is an important conclusion reached by the authors. It is now widely accepted in the system dynamics community that including all relevant stakeholders in your model building process is of utmost importance. This reflects system dynamics' assertion that action and decision making are based on deeply ingrained assumptions, known as mental models. From experience it seems clear that mental models cannot be changed based on logical advice alone. Rather, mental models can only be changed when we go through the process of learning ourselves, and feel a sense of ownership about the way in which an analysis was conducted, and the outcomes it proposes.

As such, if we want our analyses to have an impact on anyone's mental model, then they should ideally be included in the process of model building, rather than simply being shown its results. This is a valuable lesson to learn.

Reflecting on Module Five

That brings us the end of Module Five. We hope that this module has given you a brief oversight on the numerous ways in which CLDs have been used to good effect in the past, and we hope that you have taken away some valuable insights from studying these cases.

There are many, many more examples out there of CLDs being used to create change in complex systems. Below you can see links to some such cases.

This case demonstrates how the participatory modelling approach (using CLDs as the modelling language) can be used to analyse the effectiveness and potential unintended consequences of policies in large social organisations.

This case demonstrates how CLDs and participatory modelling have been used to identify feedback processes underlying ecosystem services and to foster a shared understanding of effective strategies to protect these services.

This case describes how researchers have used CLDs to explore the relationships between ecology and global finance.

Conclusion to the Course

Congratulations, you have now reached the end of the course! For a final exercise to help you reflect on what you may have learned, please go to the next page.

Image courtesy of Larry Chen from Vancouver, Canada, via Wikimedia Commons

Closing Thoughts and a Final Exercise

We hope that you have found this an interesting and inspiring course, and we hope that it will lead you to look further into the fields of systems thinking and system dynamics. Learning to think in systems and to model them is a truly iterative process. This course has introduced the basic concepts of qualitative modelling, but our experience has been that one's understanding of these basics concepts becomes more and more nuanced over time. The importance of including stakeholders in an analysis, of recognising systems archetypes, feedback loops, delays between causes and effects; these are just some of the concepts that we have found take on a new level of meaning as they are encountered in more and more contexts. We note how we have begun to see the world in a new way: to see feedback loops in everything from our own psychology to the economy, to recognising the implications of delays between causes and effects, to communicating and reflecting on our own mental models in a new way.

In this spirit of reflection, we ask that you spend a few minutes re-visiting the answers that you wrote in response to the list of questions presented at the start of Module Two. Do you think this course has had any effect on the way you might approach dealing with questions such as these? Would you now spend longer thinking about potential unintended consequences in each situation, and would you even go so far as to build a model to help you think through these issues?

Next Steps

As mentioned on the previous page, model building and systems thinking are skills that are developed with practice, and we hope that this module has given you a firm grounding to further develop those skills. There are many great resources out there that will help you develop your skills in system dynamics and systems thinking in general. Should you be interested in further pursuing your interest in system dynamics, you can find a list of courses offered on this page of the system dynamics society website. Below you can also find a list of materials that can help you further develop your skills.

Further Materials:

A course, which was also developed by Loops Consulting for The Sustainability Laboratory, uses system dynamics simulation models to increase understanding of The Lab's definition of sustainability and related 5 core principles of sustainability. It can serve as a good next step to explore quantitative system dynamics, as well as new concepts in sustainability. It can be accessed here.

This book by John Sterman serves as a great introduction to system dynamics and systems thinking (we recommend chapters 1, 3, and 5 in particular).

This website offers many enlightening articles on systems thinking and modelling.

This book by Gilbert Probst and Andrea Bassi provides a step-by-step approach to using systems thinking to solve complex problems in socio-political as well as business environment, and is a great way to explore some of the most recent applications of systems thinking and modelling in addressing pressing modern day problems.

This paper by Michael Ben-Eli serves to demonstrate how the systems perspective can bring about a more rigorous understanding of sustainability.

This website highlights some projects undertaken by a consulting firm using the systems thinking methodology. It serves as another example of how these methods can be applied in fruitful ways to help manage complex issues at the level of governments and other large organisations.

Acknowledgements

We would like to extend our sincere thanks to the Foundation for Advancement of Systems-Oriented Management Education at the University of St. Gallen, for their support in the development of this course. We would also like to thank the following individuals: Gilbert Probst, Markus Schwaninger, Andrea Bassi, Arielle Angel and the fellows of the 2018 Global Sustainability Fellowship program.